

**This Page is Inserted by IFW Indexing and Scanning
Operations and is not part of the Official Record**

BEST AVAILABLE IMAGES

Defective images within this document are accurate representations of the original documents submitted by the applicant.

Defects in the images include but are not limited to the items checked:

- BLACK BORDERS**
- IMAGE CUT OFF AT TOP, BOTTOM OR SIDES**
- FADED TEXT OR DRAWING**
- BLURRED OR ILLEGIBLE TEXT OR DRAWING**
- SKEWED/SLANTED IMAGES**
- COLOR OR BLACK AND WHITE PHOTOGRAPHS**
- GRAY SCALE DOCUMENTS**
- LINES OR MARKS ON ORIGINAL DOCUMENT**
- REFERENCE(S) OR EXHIBIT(S) SUBMITTED ARE POOR QUALITY**
- OTHER:** _____

IMAGES ARE BEST AVAILABLE COPY.

As rescanning these documents will not correct the image problems checked, please do not report these problems to the IFW Image Problem Mailbox.

RHYTHM®

CONCEPT MANUAL



i2 Technologies

RHYTHM®

Supply Chain Planner

Concept Manual

Proprietary Information of i2 Technologies, Inc.

Table of Contents

1. Reengineering the Manufacturing Planning Process	1
The Traditional Divided Process	1
Divisions Simplify the Software	3
Divisions Lead to Iteration	3
Manual Iteration is Impractical	4
Divisions Distort the Big Picture	5
The Assumptions Have Changed	6
Computing Power Has Changed	6
Planning Problem Has Changed	7
Planning Practice Has Changed	7
Planning Boundaries Have Moved	8
The Planning Process Must Change	8
Faster Tools Are Not Enough	8
Interfaced Tools Are Not Enough	9
The Process Needs Reengineering	9
Truly Integrated Planning	9
The Planner Defines the Problem	10
Integrating Master and Execution Planning	11
The Need for Propagation	11
The TIP Funnel	11
The Chicken and the Egg	14
Integrating Material and Capacity	14
Integrating the Supply Chain	15
Integrating Demand Management	15
No <i>Manual</i> Iteration	15
Integrating Manual and Automated Planning	16
<i>Truly Integrated</i>	17
2. Integrating the Expanded Manufacturing Planning Process	19
The Supply Chain	20
The Need for Visibility	21
Dependent vs. Independent Demand	22
The Traditional Approaches	23
In Practice -- Reorder-Point	23
The DRP Approach	23
DRP's Position in the Planning Process	24
Compatibility of MRP and DRP	24
Limited Usage of DRP	25
The Modern World	26
The Impact of JIT	26
Time-Based Competition	26
Modeling the TIP Supply Chain	27
Visibility	27
Point-of-Control	27
Centralized Control	27
Distributed Control	28

Hierarchical Control.....	28
Public vs. Private Information.....	29
Various Levels of Detail.....	30
Pruning the Network.....	30
Connecting the Network.....	31
Global Visibility and Local Focus	31
Driving the TIP Supply Chain.....	32
Propagating Supply and Demand.....	32
The Customer Interface	33
3. Integrating the Management of Supply and Demand	35
Promising <i>is</i> Planning	35
Truly Integrated Planning.....	36
Traditional Dis-integration	36
Visibility of Sourcing Options	36
Visibility of Cost Implications	37
Visibility of Allocation Options.....	37
Critical Optimizations May End with a Promise.....	37
TIP Must Include Demand	37
Traditional Demand Management.....	38
ATP for Quick Promising.....	38
ATP for Proper Promising.....	38
ATP Reduces Flexibility	39
Dynamic Demand Management	39
Making ATP Dynamic	39
When ATP is Not Enough.....	40
Dynamic Forecast Management.....	40
4. Reengineering Order Promising to Satisfy Business Objectives.....	43
Who's Driving?.....	43
Traditional ATP-by-Item	44
Product-Allocated ATP.....	44
ATP-by-Grade	45
ATP-by-Delivery-Lead-Time.....	45
ATP-by-Quantity	46
ATP-by-Price	46
Market-Allocated ATP.....	46
ATP-by-Customer	46
ATP-by-Market-Segment	47
Seller-Allocated ATP.....	47
ATP-by-Seller.....	48
ATP-by-Sales-Organization	48
Allocation Hierarchies.....	48
ATP-by-Product-by-Seller	49
Giving the Customer Options.....	50
5. Cooperatively Integrating Humans with Automated Planning	51
Automated Planning Algorithms.....	52
The Real Problem.....	52
Every Issue is Not Modeled	52
Planning is Rarely Automatic	53
Planning is an On-Going Problem	53
The World is Changing	53

Table of Contents

Practice Makes Perfect	54
The "Optimal" Folly	54
Cooperating with Human Planners	55
Problem-Oriented Planning	56
Automate Manual Planning.....	56
Identifying Problems	57
Allowing Infeasibility.....	57
Reasonably Near Optimal	58
Cooperating with Intelligent Local Solvers.....	58
Cooperating with Human Planners	59
Strategy-Driven Planning	59
Ask Your Teammate for Help	60
Specify the Problems to Work On	60
Specify What's 'Optimal'	61
Specify Allowed Adjustments.....	61
Specify When to Quit	61
Multi-Phase Strategies.....	61
Manual Strategies for Guidance.....	62
Special Strategies.....	62
Some Examples.....	63
MRP and then Capacity.....	63
Process Flow Scheduling (PFS).....	64
Conclusion	64

List of Figures

FIGURE 1-1: THE TRADITIONAL DIVISIONS	2
FIGURE 1-2: DISTRIBUTION PLANNING DIVISION	3
FIGURE 1-3: TRADITIONAL LEVELS	12
FIGURE 1-4: SIMPLISTIC INTEGRATION	13
FIGURE 1-5: THE TIP FUNNEL	13
FIGURE 2-1: TRADITIONAL FACTORY VIEW	21
FIGURE 2-2: THE SUPPLY CHAIN VIEW	21
FIGURE 2-3: MULTIPLE POSSIBLE FLOWS	22
FIGURE 2-4: DRP'S POSITION	24
FIGURE 2-5: UNITED MRP AND DRP	25
FIGURE 2-6: SUPPLY CHAIN PRIVACY	29
FIGURE 2-7: MODEL SIMPLIFICATION	30
FIGURE 2-8: PRUNING	31

Preface

The *RHYTHM® SCP Concept Manual* contains technical discussions about the features and usage of RHYTHM Supply Chain Planner, a Truly Integrated Planning system, which is a member of the RHYTHM® family of supply chain management products. The purpose of this manual is to introduce the many new concepts and the new planning process embodied by RHYTHM® SCP and to describe how it relates to the traditional processes.

SCP is the result of blending innovations in understanding and dealing with real customer problems together with innovations in software design and implementation that allow implementation of features that were not possible in the past. This manual is intended to address the former user-visible innovations, not the latter internal innovations.

This manual is organized into independent technical sections that each address a particular topic. The first three sections introduce a new planning process, Truly Integrated Planning. They discuss reengineering the generic planning process, without addressing SCP in particular. The following sections address SCP features that support Truly Integrated Planning, including Problem-Oriented Planning.

Each of these sections focuses on one or two concepts, describing them independently of the other concepts, and contrasting them with the traditional planning techniques. However, none of the sections in this manual bring all the concepts together to address a real problem. Instead, the *RHYTHM SCP Users Manual* illustrates how these concepts are brought together to solve the planning problems effectively.

Written by Brian M. Kennedy, CPIM.

Please forward any comments or suggestions on the presentation or content of this document to your i2 Technologies representative, or mail to the address provided.

i2 TECHNOLOGIES, INC.
909 East Las Colinas Blvd.
16th Floor
Irving, Texas 75039

1. Reengineering the Manufacturing Planning Process

The traditional planning process was based upon assumptions that were valid in the 1960s and 1970s, but no longer hold. Truly Integrated Planning (TIP) is a redesign of the overall planning process that integrates the process across several dimensions. TIP addresses the whole planning problem directly, rather than dealing with small parts of the problem separately. TIP explicitly supports the on-going, dynamic nature of planning. TIP is enabled (not merely supported) by modern computing technology. TIP software is responsible for computing the results of planning decisions and providing visibility of all the effects, leaving the human planner free to concentrate on making intelligent planning decisions.

Truly Integrated Planning is a redesign of the overall manufacturing planning process that integrates the process across several dimensions. That integration facilitates major leaps in planning capability. Such leaps that are quickly becoming necessary for all those engaged in modern, global, time-based competition. To understand why a new process is needed, it is first necessary to understand the deficiencies of the traditional process.

The current manufacturing planning process (as embodied in most MRPII products and as taught by APICS) evolved piece by piece over the last 30 years. Simple software tools were created to help solve different pieces of the planning problem. One of the first was *material requirements planning* (MRP), which computes time-phased material requirements. Other tools were then pieced around MRP to address different issues (capacity) or different levels of detail (master planning, detailed scheduling). Those tools significantly boosted planning capability and re-shaped the planning process itself.

The Traditional Divided Process

The traditional planning process is divided several times. Initially, it is divided into two major levels of detail: *master planning* and *execution planning* (see Figure 1. -1). Master planning focuses on deciding what should be built (the master production schedule) and thus what should be sold. It carries through a longer time horizon to support major, long-term planning decisions. The execution plan must have sufficient detail to drive the actual execution of the factory. However, it is only needed for a relatively short time horizon.

Each of those levels may be split again. *Master Planning* is split into:

- Production Planning
- Master Scheduling

Execution Planning is split into:

- Material Requirements Planning (MRP)
- Detailed Scheduling

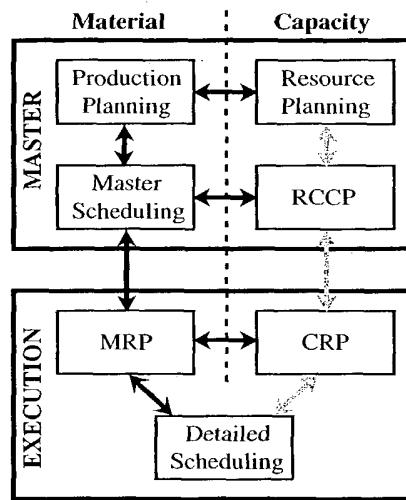


Figure 1. -1: The Traditional Divisions

The first three of those four levels are each split depending upon whether they focus on *material planning* issues or *capacity planning* issues. The net result is four groupings of seven separate process steps (and seven separate software tools):

- Production Planning and Resource Planning
- Master Scheduling and Rough-Cut Capacity Planning (RCCP)
- Material Requirements Planning (MRP) and Capacity Requirements Planning (CRP)
- Detailed Scheduling

For organizations with significant distribution networks between the customer (demand management) and the factory (master planning), some additional process steps (and separate tools) are needed. The distribution planning problem is quite similar to factory planning: demand orders for items need to be filled, various inventories (warehouses, distribution centers) of items need to be managed, and various lead times for transporting those items must be arranged. The MRP algorithm was recast to take a *Bill of Distribution* rather than a Bill of Materials, and the result is called Distribution Resource Planning (DRP).

Despite the similarity, DRP and MRP activities are not typically performed together. The interface between distribution planning and factory planning is normally the master production schedule and the demand orders (see Figure 1. -2). The distribution plan outputs demand orders that are used as input to master scheduling. Master scheduling outputs the master production schedule that specifies what to build (and thus, what can be distributed to meet the demand).

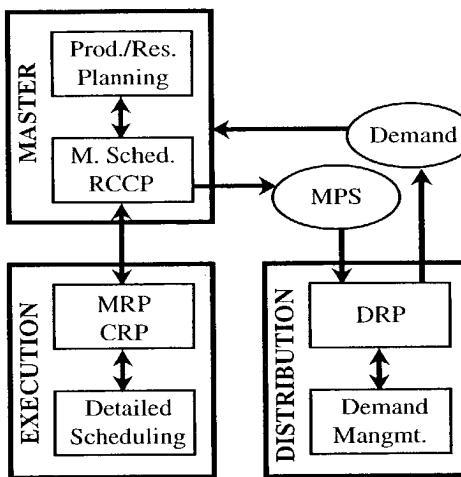


Figure 1. -2: Distribution Planning Division

Divisions Simplify the Software

At the time these divisions evolved, there were some significant practical advantages to having the problem broken up. Manufacturing planning is a highly complex problem, and it was not known how to solve the problem as a whole. Furthermore, a software tool that could address the whole problem would have required sophisticated software technology and high-performance hardware, neither of which was available. These divisions simplified the problem being addressed by each tool, and thus minimized the complexity of the software and the computing power required.

Although the divisions simplify each process step, they cannot reduce the inherent complexity of planning a factory (or supply chain) well. Rather, the divisions simplify the software by shifting the burden to the human planner.

Divisions Lead to Iteration

At the higher planning levels, a longer time horizon is considered but with much less detail. At the lower levels, much more detail is considered but with only a short time horizon. The material tools largely ignore capacity constraints, and the capacity tools largely ignore material constraints. The distribution tools ignore factory details, and the factory tools ignore details of the demand distribution.

Since each module only considers part of the whole problem, the computational requirements are much smaller. However, undesirable

decisions made with one tool may not be detected until the results are propagated to another tool. Thus, iteration between the various tools is an essential part of the traditional planning process. The whole manufacturing planning problem typically cannot be solved well without iterating between these various tools.

For instance, the many decisions made in master scheduling have implications on capacity usage. Those implications are not known until a master schedule is created and run through RCCP. If RCCP determines that a master schedule is not feasible with the attainable capacity, the human planners must return to the master scheduling module and revise the master schedule as indicated by the RCCP results. RCCP again evaluates the new master schedule and any new problems require an additional iteration.

Once the master schedule "passes" RCCP, it is passed on to MRP and CRP for more detailed planning. If the more detailed planning determines problems in the master schedule, the human planner must return to the master scheduling and RCCP modules to again revise the master schedule. The same potential need for iteration back to master scheduling exists through CRP and detailed scheduling.

Furthermore, decisions made in the factory plan may have significant consequences in the distribution plan and vice-versa. Such decisions require iteration through distribution planning back into factory planning, which includes all the iterations described previously.

Manual Iteration is Impractical

A human planner *can* move back and forth between the modules of the traditional planning process and develop a good, feasible plan. At least, *theoretically* it can be done.

Practically, however, such manual iteration requires significant human effort. The planner must leave the current tool in such a way that the results are available to the other tools. Each of the other tools must be started and must read in the new results. The human planner, knowing the changes made in the previous tool, must then begin the proper analyses for each tool to display the interesting effects of those changes. If adjustments must be made, the results of that tool must be saved. As this continues, the collection of all unanalyzed changes must be remembered so that the planner can properly analyze the various tools to display all possible effects. Of course, knowing what the interesting effects of a change could be requires considerable knowledge. The alternative to mentally keeping track of the effects to be investigated is to analyze all the results (which is no small task).

Given the manual effort involved, it is not too surprising that the great majority of real-world planning cycles involve *zero feedback* iterations (that is, they involve one pass straight through the planning tools). It is quite rare that more than one feedback iteration is performed. To avoid needing to iterate, human planners are often overly conservative in the earlier tools to prevent the later tools (capacity or more detailed) from detecting problems that would require iteration. To avoid iteration between tools, they cleverly make higher level decisions that are least likely to cause problems in the later tools. This may be done instead of making decisions that are most profitable to the company.

Once in the lower-level or later tools, planners diligently work around poor decisions made in previous tools rather than go back and change those decisions. In fact, many organizations set up operational constraints that do not allow changing the higher level plans within the horizons of the more detailed tools.

For example, a material-constrained firm will do material planning first. They use assumptions and planning rules that avoid reaching any capacity constraints. They may also purchase excess capacity to keep use of materials low or use planning rules that keep use of materials fairly constant. In that way, the capacity tools rarely detect an overload.

Likewise, a capacity-constrained firm will do capacity planning first. They use assumptions and rules that avoid reaching any material constraints. This typically involves decisions that enforce either stable material usage or higher inventory levels (which effectively relax the material requirements).

Most organizations' plans, despite being conservative overall, are infeasible by the time they are being executed. As a result, informal mechanisms, such as emergency expediting, emerge to work around the formal mechanism. Though necessary given infeasible plans, emergency expediting typically reduces the effective capacity of the organization. This results in a situation that requires yet more expediting.

Traditional planning practices involve the human effort to manually control iterations through the separate modules, the considerable thought and manual effort in guessing the effects of decisions across modules, and the need to work around those effects later. All such efforts are a waste. Controlling iterations between tools and propagating the results of decisions is work ideally suited for computers.

Many of the monthly and weekly planning cycles of the past are simply inappropriate today, as are many of the activities that once occupied a great deal of a human planner's time. Such routine activities should be off-loaded to a software system.

Divisions Distort the Big Picture

Not only are such iterations between tools cumbersome, they also make it difficult to see the big picture. The planner can only see the effects being modeled by the software tool in use. The many other effects must either be estimated mentally or computed by switching to each of the other tools.

The human planners' responsibility is to create and maintain a good plan for the whole manufacturing organization. Even if no software tool is considering the whole plan the human planner must.

Peter Senge begins his book *The Fifth Discipline* with this observation:

From a very early age, we are taught to break apart problems, to fragment the world. This apparently makes complex tasks and subjects more manageable, but we pay a hidden enormous price. We can no longer see the consequences of our actions; we lose our intrinsic sense of connection to a larger whole. When we try to "see the big picture," we try to reassemble the fragments in our minds, to list and organize

all the pieces. [...] the task is futile [...] Thus, after a while we give up trying to see the whole altogether.¹

Dividing the problem indeed simplifies the software tools that deal with each piece. However, the burden of understanding and managing the whole is shifted to the human planner. No matter how effective we make the software for those pieces, the effectiveness of the overall planning process will be limited by the human planners' ability to understand and mentally manage the whole.

Managing a great deal of information, organizing it, and presenting it from various vantage points are things that computers do well. It is better to task the computer with managing the whole plan and with providing the human planner full visibility. That leaves the human planner free to intelligently plan and deal with future problems.

The Assumptions Have Changed

The design of the traditional divided planning process (and the software tools that embody it) was based upon a number of assumptions that have since changed dramatically. Computing power and the basic understanding of the manufacturing planning process have advanced tremendously. Further, the planning problem itself has also been moving as the marketplace has globalized and reengineered itself. Many of the changed assumptions call for rethinking the basic planning process.

Computing Power Has Changed

The APICS' CPIM Systems and Technologies Certification Review Course begins discussion of implementations with the observation,

“All implementations, whether of new systems or new technologies, should represent process and performance improvements for the firm. This, however, is not always the case. Implementations often simply represent mechanization of the current process. Though this will produce some benefits, there are much more significant benefits to be gained if the process is reengineered at the same time that systems or technologies are being applied.”

Since the traditional APICS planning process was conceived in the 1960s and 1970s, computing power has doubled many times over. At the same time software design has advanced significantly. Advances in planning and scheduling technology coupled with advances in optimization algorithms let software take care of most of the routine decisions that must be made. Advances in graphical user interfaces let interactive software efficiently communicate information and true understanding to human planners.

Most of that technology has been applied to simply enhance the support for the traditional planning process. Far more can be accomplished by revising the process to take advantage of the new technological capabilities. For example, more visibility can be provided and more analyses can be

¹Senge, Peter M., *The Fifth Discipline: The Art and Practice of the Learning Organization*, Doubleday, NY, 1990, p. 3.

accomplished within times that are comfortable for an interactive user. Moving from batch processes to interactive software completely changes the nature of the tools that are possible.

With hundreds of megabytes of computer memory available on a desktop, modeling all the details of a supply chain is perfectly reasonable. Unless trends change dramatically, we can expect ten times as much computer capacity in just a few years. Global visibility of a supply chain from interactive software should now be expected. And once one competitor has global visibility, the others soon need it.

Planning Problem Has Changed

The competitive driver of the 1980s was quality and reliability. The concepts of *Total Quality Management* (TQM) and *Just-In-Time* (JIT) greatly transformed the way manufacturers do business and thus changed the planning problem. Excess inventories, which in the past had been considered a necessary part of manufacturing practice, were no longer tolerated. The slack in the systems was eliminated. Without that slack, the need for either greater stability and consistency or more accurate and more dynamic planning capabilities increased dramatically.

The competitive drivers of the 1990s are time and service (responsiveness and flexibility). Concepts such as Agile Manufacturing, Time-based Competition, and Supply-Chain Management have begun to significantly impact the way manufacturers compete. The notions of stabilizing flow and reducing variation are no longer competitive practices. Instant, reliable, and aggressive order quotation is becoming standard practice as markets gain global visibility. That leaves more accurate and dynamic planning capabilities as the only competitive option.

Planning Practice Has Changed

The concepts of *Synchronous Manufacturing* and the *Theory of Constraints* (TOC) have changed the planning activity itself. The fundamental understanding of the manufacturing process, its behavior, and how it is controlled has advanced tremendously in the last 20 years.

Eli Goldratt and Bob Fox conclude their book *The Race* with the motto, "The sum of the local optimums is not equal to the global optimum."² Although this principle is now generally accepted, the current divided planning process cannot address the global manufacturing problem well. To solve the global planning problem, the global effects of planning decisions must be presented. Presenting only local effects is of little benefit. Goldratt and Fox argue, "Schedules should be established by looking at all the constraints simultaneously."³

The problem as a whole must be directly addressed. The many different issues involved in planning manufacturing operations are highly interrelated. Addressing these issues separately can indeed simplify each problem being considered, but a separation actually makes it more difficult to develop a

²Goldratt, E. and Fox, B., *The Race*, North River Press, NY, 1986, p. 179.

³Goldratt, E. and Fox, B., *The Race*, North River Press, NY, 1986, p. 179.

good plan that considers all the relevant issues. Without direct visibility of the global results of a planning decision, it is very difficult to make decisions that are good from a global perspective.

Given global visibility of the planning problems that will occur in the future, planners must have the time and the tools to resolve those problems long before they even occur.

Planning Boundaries Have Moved

The world in which manufacturers compete has changed dramatically since the traditional planning process was originally developed. In particular, the global relationships between a manufacturer and its suppliers and customers have changed tremendously. Supply chain management has significantly extended the boundaries of the planning problem. JIT concepts have changed the supplier-consumer relationship at each link in the supply chain.

Manufacturers can no longer compete successfully with blinders on. To respond faster than their competitors they must look out beyond their own organizations to the supply chain upstream and downstream.

Planning without visibility of the whole supply chain is suboptimal and will soon become unacceptable. If manufacturers are to compete successfully in the global marketplace, they must provide responsiveness and flexibility not only to their direct customers but to their entire supply chain downstream.

The new computer communication technologies have been applied to better connect these organizations. But the overall approach to solving the planning problem has not changed in response to the change in the problem and the implementation options.

The Planning Process Must Change

Many have been calling for improvements to the tools that support the planning process, for faster and more interactive tools, and for interfacing the tools to make them "seamless." But is that enough?

Faster Tools Are Not Enough

The need for iteration has resulted in a push for software tools that are significantly faster than before. Such speed makes manual iteration more viable. Instead of needing several hours to perform an iteration back through master scheduling, new computing power allows iteration in minutes.

Such speedups are a significant step forward. They allow the number of feedback iterations to grow from zero or one to several. That lets human planners produce significantly better plans, particularly when they compare them to global criteria.

However, human planners are still forced to manage the iteration through the tools. Immediate global visibility of the effects of decisions remains unavailable, so human planners must still trade-off effort. They must estimate what the other tools would tell them versus the effort of actually using those tools and running the appropriate analyses. The amount of iteration continues to be limited to a relatively small number due to the human effort involved.

(Performing hundreds of iterations remains unreasonable; thousands are impossible.)

Interfaced Tools Are Not Enough

Many manufacturing software vendors and consultants have been stressing the importance and value of “integrating” the various software modules. Books have been written about “integrated” planning. But all that they are talking about is making them work together, making them able to pass information and draw upon the same database. Such integration, better termed “interfacing”, merely *allows* iteration to be performed. That is a start but not nearly enough.

The interfaced tools cannot cooperate to form a single picture. Analyses involving all the information is not possible. Global visibility can only be obtained by moving among the tools and mentally collecting and analyzing the information.

The Process Needs Reengineering

The basic planning process (and the software module divisions) that evolved through the 1970s has remained in place. This is true despite the fact that the assumptions on which the divisions were made no longer hold and in fact, have changed dramatically. The traditional manufacturing planning process does not meet certain criteria that successful manufacturers of the future must. Some problems with the traditional manufacturing planning process are:

- It does not take full advantage of the computing power available today.
- It does not properly address the planning problem.
- It is not well suited for many modern manufacturing techniques.
- It does not address the whole supply chain.

The process can be further refined and automated, but these problems will not be resolved without reengineering the process itself. *Truly Integrated Planning* (TIP) is such a redesign.

Truly Integrated Planning

The phrase “truly integrated” in TIP calls for bringing the independent modules into *one tool*. This tool can immediately pass on the effects of planning decisions and begin the appropriate analyses. It involves providing immediate visibility of the important effects of the planning decisions that are made. It requires getting humans out of the computational iteration chores, so they can concentrate on making the more important intellectual decisions and trade-offs.

Such integration of the tools enables integration of the *process* itself. The iterations occur within the software so there are no process divisions to iterate through. TIP requires that the planning software automatically perform the appropriate computations for each planning decision made, and provide immediate visibility of all effects to the human planner.

Computers are very good at iterating through tasks until the answer is found. They are also very good at starting other subprograms to determine and pass on the effects of decisions. Let computers do what computers do well. Leave the human planners free to use global visibility to make the hard decisions and devise innovative solutions.

Based on a significantly different set of assumptions, Truly Integrated Planning (TIP) involves many fundamental changes from the traditional planning process. The tools and process steps are *integrated* into one. The planning problem is addressed as an *integrated* whole allowing for global optimization. Material and capacity issues are *integrated*, detailed near-term issues and longer-term issues are *integrated*, and factory and distribution issues are *integrated*. An *integrated* picture of the problem provides instant and global visibility of the effects of planning decisions throughout the supply chain. The on-going planning efforts are *integrated* over time to effectively evolve better plans. Automated and manual planning are cooperatively *integrated* so that they can each be used to their maximum potential. That is the level of *integration* called for by *Truly Integrated Planning*. Separate tools that have simply been interfaced together do not satisfy TIP requirements.

The software planning system is an integral part of the definition of TIP. Such dependence is not unprecedented or even new. Gopal and Cypress state, “Information technology (IT) is the thread running through the enterprise – the link that truly integrates manufacturing distribution and customer delivery. Its innovative use defines the competitive arena, and its role in integrated distributions is that of key enabler. In other words, it is an integral part of, not a support for, the process.”⁴ TIP, as a reengineering of the planning process, directly incorporates the new role of software as key enabler. TIP does not leave software as an “automation” afterthought.

TIP requires integration across multiple dimensions of the planning process, and integration requires further integration. For instance, requiring software to address the entire supply chain further requires the software to integrate the capabilities needed by the various types of organizations involved. Repetitive, batch, process, discrete, and rate-based organizations commonly occur in the same supply chain, or even the same organization. TIP software must cover all types of problems to provide global visibility of most supply chains.

The Planner Defines the Problem

The TIP directive to “address the whole problem” does not mean that every detail about a manufacturer must be modeled. The “whole problem” is everything that the planner cares about when making planning decisions. For example, most of the machines in a typical factory are non-bottlenecks. There is no need to model such machines in great detail. In fact, there may be no need to model them at all. Some planners may not explicitly plan labor and thus it need not be modeled. Much of the surrounding supply chain is too weakly affected by local planning decisions to be of concern to the planner.

⁴Gopal and Cypress, *Integrated Distribution Management*, Business One Irwin, 1993, p. 145.

TIP software does not *require* the planner to model any more than is required by a typical MRPII system. Rather, TIP software is required to be *able* to model anything in the supply chain. The typical planners care deeply about some of the issues in the following discussion, but they may care very little about others.

Integrating Master and Execution Planning

The split between master and execution planning is not artificial. They have different uses, horizons, and levels of detail. Integrating them does not mean that they become one. It means they are implemented in one software package that can automatically and seamlessly move between them.

It is wasteful to ignore near-term detailed plans when making longer term plans. Near-term planning decisions may have important effects on both the near-term and longer-term plans. These effects are only visible if the plans are integrated.

The Need for Propagation

To illustrate the importance of propagating planning decisions between master and execution levels, consider the master-level decision to build twice as many "widgets" next month. The decision may require that building some of the components starts in the next few days. The near-term plans need to see that demand. In addition, the master-level decision should take advantage of the detailed near-term plans when determining feasibility.

Similarly, if it is decided to break a setup on a machine in the near-term horizon to satisfy an expedited order, it may affect whether all those widgets can still be built. Similarly, if a machine goes down for the next week, it may not be possible to build all those widgets. The long-term plan needs to see the different conditions.

If the delay of next months' widgets due to breaking that setup is accepted, many of the other widget components are not needed this week. They can be delayed until next week allowing for more urgent orders to be performed first.

Note the effect: the near-term planning of the down machine or broken setup causes a long-term delay, and that delay changes the priorities on other near-term jobs. Without the integrated long-term plan, the effect of one near-term decision on other near-term plans may not be visible. To see the effects with the traditional manufacturing process, you must manually move between the master planning and execution planning tools. Thus, integration of master and execution planning is necessary to provide the needed visibility.

The TIP Funnel

Integration requires a single model of both the near-term and long-term plans. However, it is futile to make detailed decisions for long-term plans because the detail will surely be wrong (most environments are far too dynamic). Any effort expended to compute detailed plans for longer horizons is a waste.

The near-term plans should be detailed. The long-term plans, though in the same model, should be computed in much less detail. Near-term decisions

should be propagated through the long-term plan, long-term planning calculations should take advantage of the near-term details, and excess useless detail should not be computed for long-term plans. The "TIP Funnel" is a useful conceptual analogy for the integration of master and execution planning.

The traditional planning model is depicted in Figure 1. -3. The goal of production planning is to create a long-term plan that is very coarse in detail. The goal of master scheduling is to form a more detailed (but still aggregated) plan for a shorter horizon. Similarly, the goal of MRP/CRP is to further increase the planning detail for an even shorter horizon. And finally, scheduling plans in great detail but for only a very short horizon.

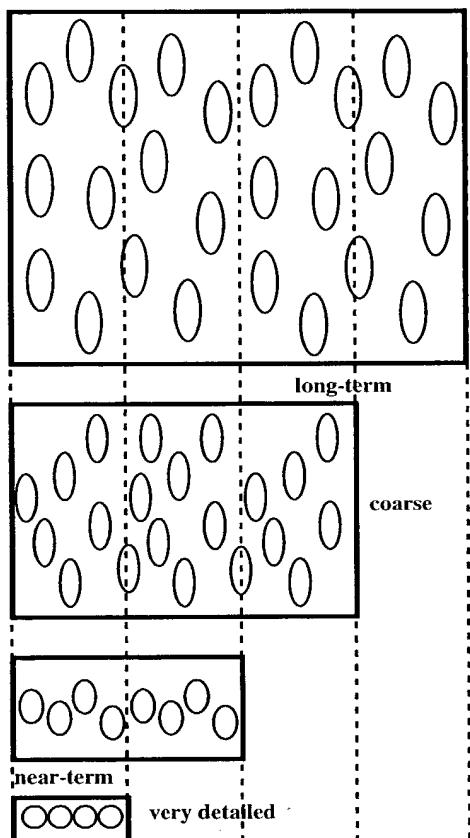


Figure 1. -3: Traditional Levels

A simplistic attempt at integrating these levels would involve interfacing the plans indirectly by using the output from the more detailed plans to impose coarse constraints on the less detailed plans. The longer horizon results can be used as demand input to the near-term plans. This is represented graphically in Figure 1. -4.

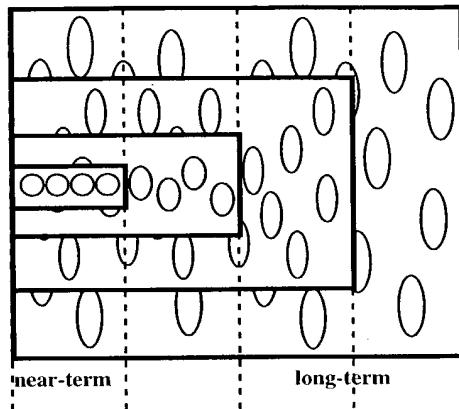


Figure 1. -4: Simplistic Integration

As time passes and the on-going planning process continues, long-term plans move into nearer horizons increasing the level of detail. The plans eventually move into the nearest and most detailed horizon. Thus, plans are continually revised and refined to greater levels of detail as their time nears.

Rather than overlaying disjoint plans with different detail levels and horizons, more effective integration can be obtained by merging the plans into one. That one plan, in contrast, can vary in detail with the horizon. This is depicted by the "TIP Funnel" (see Figure 1. -5). As time passes the funnel gradually slides forward increasing the detail of the plans.

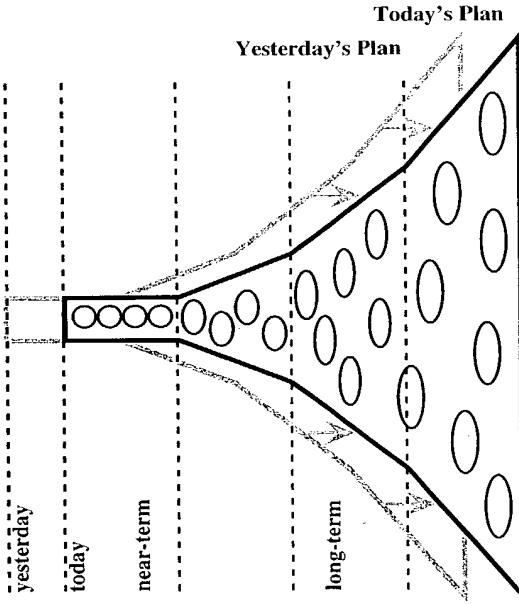


Figure 1. -5: The TIP Funnel

When an activity first enters the planning funnel (i.e., enters the production planning horizon), it is planned very coarsely. What actually occurs on dates far in the future may vary widely from the original plans. As the funnel moves forward (time passes), the plans are refined narrowing how much they vary from what actually occurs. Eventually, the plans reach the neck of the funnel where they are made in great detail. It is hoped that these plans accurately reflect what happens.

The Chicken and the Egg

The traditional planning process is often presented as a top-down procedure. The production plan is developed first and then refined into the master schedule. The master schedule is then broken out into individual items in MRP. The procedure continues in this manner. Vollmann, Berry, and Whybark observe, "An interesting chicken-and-egg question sometimes arises about the production plan and detailed plans that result from the MPC system. Conceptually, production planning should precede and direct MPC decision making. In some firms, [...] the first production plans are no more than a summation of the individual detailed plans. They're the result of other detailed decisions, rather than an input to those decisions."⁵

The conceptual problem results from failure to recognize that planning is *not* the process of building a plan. Rather, it is the on-going process of maintaining a plan. It is wasteful to throw out the detailed plans created yesterday, perform high-level planning, and then completely redo detailed planning based on that high-level plan. Instead, the plans should evolve as time passes.

Thus, the production plan is *both* an aggregated view of the detailed plans and a higher-level plan from which executive decisions can be made. Top-down, bottom-up, and middle-out are all valid directions for propagating planning decisions in an on-going process. With one integrated plan, working at any level immediately results in propagation to all levels.

Integrating Material and Capacity

It is widely accepted that both material and capacity plans must be developed and be consistent if either plan is to be feasible. Vollmann, Berry, and Whybark state, "Capacity plans must be developed concurrently with material plans if the material plans are to be realized."⁶ However, they do not propose altering the traditional planning process to bring these two activities into one.

It is no longer competitively acceptable to operate with either excess capacity or excess inventory simply to work around weak planning capabilities. As companies push towards lower inventory levels and lower levels of excess capacity, the planning problems in each area increase. As material and capacity constraints begin to interact, tools that address both concerns simultaneously become essential.

⁵ Vollmann, Berry, Whybark, *Manufacturing Planning and Control Systems*, Third Ed., Business One Irwin, 1992, p. 277.

⁶ Vollmann, Berry, Whybark, *Manufacturing Planning and Control Systems*, Third Ed., Business One Irwin, 1992, p. 154.

Addressing both concerns together requires significantly more sophisticated software, but it does not require anything revolutionary. In fact, some of the better "finite-capacity scheduling" tools have the basic ingredients necessary to address both the material and capacity problems together.

Integrating the Supply Chain

The preceding paragraphs describe the TIP process for the traditional manufacturing problem. However, for many manufacturers, addressing the whole problem includes expanding the traditional boundaries of the planning process to include the surrounding supply chain. In the modern world of global competition and tight cooperation between supplying and consuming manufacturers, along with the advent of time-based competition, seeing the immediate impact of decisions on the supply chain is essential.

Integrating the entire supply chain involves much more than just applying the traditional planning process to the bigger problem. In fact, it involves more than just applying the TIP-reengineered planning process described in this section. It involves further reengineering of the planning process. For a software tool to address the entire supply chain, it must integrate across industry boundaries (e.g., steel, automotive), manufacturing boundaries (e.g., batch, repetitive; make-to-order, make-to-stock), and organizational boundaries (e.g., manufacturing, distribution, sales, retail, different companies, etc.). It must address additional issues such as authority, control, privacy, distributed modeling, and cooperative planning.

Integrated planning of the full supply chain with TIP is discussed further in the section, "Integrating the Expanded Manufacturing Planning Process".

Integrating Demand Management

How quickly can you get what a customer wants, the way they want it, when they want it? That is the competitive question of the 1990s, and an answer is often expected while still on the phone. Further, the reliability of the answer must be very high.

Without sophisticated planning capabilities integrated with order promising, the necessary speed and reliability of the answer are not possible. Reliable answers can be pre-computed but flexibility is lost. As a result, potential sales are lost.

Supply chain orientation and JIT relationships have also led to more complex agreements that effectively constrain the planning process. Marketing and sales efforts can take significant time to ramp-up effectively. That, too, can be a critical bottleneck that must be understood during planning.

Full discussion of truly integrated demand management and order promising is included in the sections "Integrating the Management of Supply and Demand" and "Reengineering Order Promising to Satisfy Business Objectives".

No Manual Iteration

The phrase "truly integrated" in TIP means far more than working seamlessly together. It means immediate global visibility and elimination of all *manual*

iteration. When a human planner makes a decision at any level of detail, the computer should be responsible for immediately propagating the results of that decision. It should occur through both material and capacity analyses at all relevant levels of detail and by reporting any problems. The human planner should not be involved in computational propagation and bookkeeping.

The effects of such a fundamental design change are dramatic. Feedback from different analyses becomes fast and painless because no laborious human tasks are involved. Iteration is done as needed (typically, the more the better). All parts of the plan become responsive to each other.

None of the effects of excessively conservative planning are seen. Neither excess capacity (low material use), nor excess inventory is created. The need for multiple manual iterations between capacity and material planning efforts is avoided. Similarly, restrictive organizational and planning rules can often be relaxed, allowing planners much more flexibility in dealing with their ever-changing environment.

Although over-conservatism is reduced, plan feasibility is still increased. The effects of every decision can be quickly propagated throughout the plans. With the added flexibility provided by elimination of restrictions, planners can respond to changes intelligently. Trying to plan sufficient slack into the system, so that expeditors can respond to the changes without completely invalidating the plans, is no longer necessary.

For the first time, global optimization of the firms' plans becomes possible. With the computer performing all such iteration immediately, the global effects of planning decisions become quickly visible.

Integrating Manual and Automated Planning

Automated planning algorithms can handle the majority of the simpler planning decisions and edits. Human planners are free to concentrate on the difficult problems, major issues, and trade-offs that require added knowledge of the operations.

It is unreasonable to assume that all information can be made available to the planning software. There are always more possibilities in the real world than the software knows about. Most of the hard decisions and trade-offs must be made by the human planners. On the other hand, manufacturing planning in the real world is a huge task. It could easily consume an army of human planners unless the majority of it is computer automated.

Planning is an on-going activity. It should be a long-term integrated effort, not a series of disjoint repetitions. Planning is *not* the process of building a plan. Rather, it is the on-going process of maintaining a plan. In the real world, the planning problem rarely involves computing a plan from scratch. There is always an existing plan in place into which much planning effort has been poured. It makes little sense to throw out that plan and build a new one.

Given the on-going nature of planning, it is critical that automated and manual planning be freely intermixable. It is not acceptable for automated algorithms to ignore (and step on) the existing manual edits, although most automated algorithms do just that. The automated algorithms must take the existing plan, along with the edits performed upon it, as an additional basic

input. The algorithms should be designed to continue improving the existing plan, primarily adjusting for things that have changed, without destroying the manual edits on which the human planner worked so hard.

Automated algorithms can be extremely useful during and after manual editing. Manual editing can generate more new problems than it fixed, but the new problems may be easier for the computer to solve. However, many algorithms simply cannot be applied in such a manner. Designing an algorithm that can be used cooperatively with manual editing fundamentally affects the nature of the algorithm.

Many of the analyses and computations used by the automated algorithms can be highly useful to human planners as they make manual decisions. For example, consider the analyses that must be done by the automated planner to choose an alternate resource or routing. It must find the places where an alternate may be chosen and then evaluate the potential gains from selecting each alternate. A human planner also likes to see the list of alternates that are possible and the evaluation of the alternates that are most promising to explore.

The manual tools should be able to work with those same analyses. Doing so requires that the manual tools and the automated algorithms operate in a consistent and complementary manner.

Truly Integrated

In summary, “*truly integrated*” involves integration across multiple dimensions. Material and capacity planning are integrated into one problem. Production and distribution planning are integrated, and the entire supply chain can be addressed as one problem if desired. To support this integration, different industry types, manufacturing types, and organizational types of planning software are integrated into one. Planning at various levels of detail, with different horizons, is integrated into a single “funnel” with detail varying over the horizon. Successive planning efforts in the on-going planning process are integrated such that yesterday’s results are used in today’s planning effort. The automated algorithms and manual planning tools are fully integrated so that they can take advantage of each other.

Truly Integrated Planning is a reengineered planning process enabled by advances in hardware, software, optimization, and planning technologies. Without such reengineering of the planning process, human planners continue to be limited by their ability to mentally assemble the big picture of the plan from their efforts in the separate process steps. The automated planning tools remain limited to addressing only part of the global problem.

2. Integrating the Expanded Manufacturing Planning Process

The Truly Integrated Planning (TIP) process directly addresses the planning problem as a whole. For many organizations, the whole planning problem involves a larger supply chain, and thus visibility of the supply chain is necessary. Modeling the larger supply chain raises numerous integration issues that must be handled by the enabling software system, including issues of authority, privacy, levels of detail, interconnection, organizational types, manufacturing types, and cooperative control.

The previous section "Reengineering the Manufacturing Planning Process" introduces a new planning process. This section continues the definition of truly integrated planning, focusing on its application to the supply chain planning problem as a whole.

Based on a significantly different set of assumptions, Truly Integrated Planning (TIP) involves many fundamental changes from the traditional planning process. The tools and process steps are *integrated* into one. The planning problem is addressed as an *integrated* whole allowing for global optimization. Material and capacity issues are *integrated*. Detailed near-term issues and longer-term issues are *integrated*. Factory and distribution issues are *integrated*. An *integrated* picture of the problem provides instant and global visibility of the effects of planning decisions throughout the supply chain. The ongoing planning efforts are *integrated* over time to effectively evolve better plans. Automated and manual planning are cooperatively *integrated* so that they can each be used to their maximum potential. Cooperative *integration* for maximum potential is the level of *integration* called for by *Truly Integrated* Planning. Separate tools that have simply been interfaced together do not satisfy TIP requirements.

In today's global marketplace, where responsiveness and time are the major market differentiators, managing the flow through the entire supply chain to the final customer has continued to grow in importance. The demands made on a manufacturer by the client are dependent upon demands made by the client's customers. This dependence is a natural progression. Each manufacturer, distributor, and sales organization is dependent upon the same end customer demand. All of them are more successful if they cooperate better in satisfying that demand. If they do not successfully supply the demand, another supply chain will.

For organizations whose performance is dependent upon a larger supply chain, the principles of TIP call for directly addressing the entire supply chain in the planning process. Most large organizations have many possible alternative flows through their supply chain. Some of the most complex planning decisions are in choosing a particular flow path to satisfy each demand. Effectively planning such a network requires immediate visibility of all possible flows. Efficient operation of such a network cannot be achieved by planning small individual pieces of the larger problem.

Andre Martin observes, "We must stop developing stand-alone systems, and begin looking for linkages which provide visibility from one level to another. [...] Remember, any change at any point will be felt elsewhere in the [supply chain] -- it's only a matter of time and quantity before the shock wave hits. If we aspire, as we should, to fully synchronize [the supply chain], we must begin to install planning and scheduling systems, and then link them together."⁷

TIP principles further argue that those linked planning systems should provide immediate visibility of the effects of local decisions on that larger supply chain. It is not enough to have slow propagation of the decisions made at each organization. The planners at each organization should receive immediate feedback on the probable effects of each planning decision that they make.

Applying the *Truly Integrated* Planning process to the entire supply chain involves much more than simply expanding the model. Issues of consistency, privacy, authority, and focus make the supply chain problem significantly more difficult than just building a larger model. The supply chain is made up of more than just manufacturers. There are warehouses, distributors, wholesalers, retailers, sales professionals, and order entry and customer service personnel. Transportation scheduling must take into account many issues that are not present in factory operation scheduling. All of these different needs and personnel must be integrated in a mutually productive manner. They must be integrated without sacrificing local efficiency or violating privacy and authority issues.

The Supply Chain

Typically, manufacturing planning software takes in a model of the factory, status information on its current state, and demand information on what needs to be satisfied. Its job, along with the human planners, is to plan the activity of the factory such that the demand is satisfied efficiently and effectively (least cost, most profit, least inventory, highest service level, etc.).

Fundamentally however, the manufacturer's job is not to satisfy the incoming demand orders. It is to satisfy the end customers. It does not really matter if ABC Manufacturing satisfies or misses the due dates on its incoming orders. It does not matter if their finished goods inventory stocks out or if their distribution center stocks out. What matters is that their product never stocks out on retail store shelves (assuming ABC is a make-to-stock manufacturer of consumer products). Running out of input materials from

⁷ Andre Martin, *Distribution Resource Planning*, Revised Ed., Oliver Wight Pub., VT, 1993, p. 40.

their suppliers is no excuse. ABC's real planning problem involves all the members involved in keeping the retailers' shelves stocked: the retailer, the distribution centers, the warehouses, the factories, and the suppliers.

This path, from raw materials through finished product and into customers' hands, has been termed the "supply chain", "supply-demand pipeline", "market channel", "distribution network", "virtual corporation", and so on. Managing that supply chain in an optimal way has taken on elevated importance as flexibility and responsiveness have become key differentiators in today's increasingly global marketplace.

The Need for Visibility

Figure 2. -1 depicts a factory using the traditional planning process. It builds two products, "widgets" and "gadgets". For both products, the factory receives demand orders for this month that are well over forecast. The factory does not have capacity to prevent stock-out on both, so the factory workers work hard to minimize the stock-outs of both products.

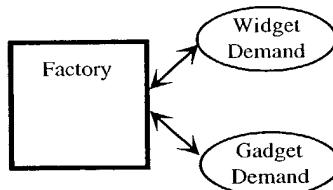


Figure 2. -1: Traditional Factory View

In reality, however, only one of the products, the widgets, is truly selling over forecast. The gadgets just happened to hit reorder points for various distributors at the same time (see Figure 2. -2). Thus, there are actually plenty of gadgets in the system to continue to satisfy demand. The cumulative inventory of widgets, on the other hand, is truly being depleted. But the factory could not tell the difference because the workers only see the demand orders. If they instead could see all of the inventory between them and the end customers, then they would be able to increase fabrication of only widgets. The factory would have sufficient capacity to match the demand.

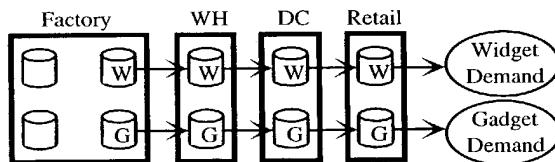


Figure 2. -2: The Supply Chain View

When using the traditional planning tools each stage cannot see the others, so they must each forecast (i.e., guess) what they believe the other organizations will do. They must each carry their own safety buffers to cover variations

from their guesses. Each member of the supply chain maintains buffers for fluctuations in supply from upstream and buffers for fluctuations in demand from downstream. All the upstream and downstream members are doing the same. The net effect is that the total inventory in this system is much higher than necessary. High inventory has a direct effect on the products' competitiveness in the market just as if the supply chain was one firm with those same excesses. As a result, all companies involved in the supply chain suffer. (If the end item is less marketable, then the demand on every company in the chain is reduced.) Excess inventory means excess investment, excess lead time, lower responsiveness, more waste, and so on.

This discussion has revolved around a simple linear supply chain. However, many supply chains consist of a network of possible flows (see Figure 2. -3). If one plant lacks the capacity to satisfy the demand, a sister plant may be able to cover it. The cost may be somewhat higher to use the sister plant, but that cost may be much less than the cost of a stockout or late order. Planning which flow(s) to use through the supply chain becomes a critical issue that requires seeing each possible flow.

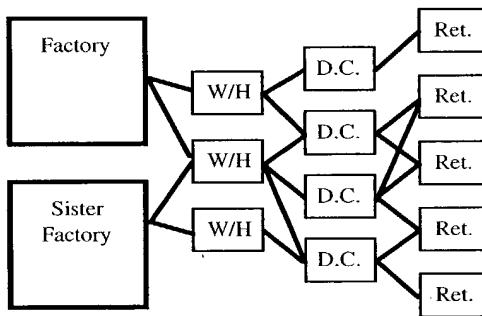


Figure 2. -3: Multiple Possible Flows

Dependent vs. Independent Demand

If the supply chain pictured in Figure 2. -2 was one large company, then the duplicated inventories based upon the multiple independent forecasts would probably not be tolerated. Such waste would be eliminated by taking advantage of *dependent* demand. The downstream members pass demand to upstream members. If a plan has been developed for the downstream member, then the demand on the upstream members is known. As stated in the APICS Dictionary definition of dependent demand, "Such demands are therefore calculated and need not and should not be forecasted."⁸

The only forecasts that should be generated (guessed) are for *independent* demand, demand that comes from sources that are not being planned within the model. Forecasting dependent demand is simply bad practice.

The logic of MRP was founded upon this principle. MRP is simply an engine for propagating dependent demand upstream. And that demand must be time-phased appropriately.

⁸Cox, James F., John H. Blackstone Jr., and Michael S. Spenser, Ed., *APICS Dictionary*, Eighth Ed., 1995, p22.

The primary purpose of integrating supply chain planning efforts is to propagate planning decisions throughout the supply chain, so that the effects are visible to all. Doing so eliminates the need to forecast (guess) what the other members are going to do. Integrated planning eliminates the need for the many safety buffers created to compensate for poor guesses.

The Traditional Approaches

Supply chains exist whether we plan them or not. Although the importance of such planning is generally accepted, and although better mechanisms are known, common practice involves simplistic planning mechanisms that do not adequately account for the surrounding supply chain. Before proposing superior mechanisms, it is important to understand the existing mechanisms.

In Practice -- Reorder-Point

Reorder-point is a simplistic computation designed to statistically characterize both upstream and downstream behaviors. It involves encoding the forecast of upstream and downstream behavior into a pair of numeric values: the *inventory quantity* at which an order should be placed and a *quantity to order*. The many derivative forms of reorder-point change the number and nature of those numeric values, but the overall concept is the same.

Such over-simplification of the problem can lead to truly misguided planning decisions. Despite that, as Andre Martin observes, "Today, in fact, more than 99 percent of all retailers, wholesalers, and distributors use the classic reorder-point technique and its derivatives to manage and control inventories and conduct purchasing."⁹ Martin goes on to provide detailed examples of how reorder-point is fundamentally flawed. He (like many others) argues that Distribution Resource Planning (DRP) can provide tremendous performance improvements over reorder-point techniques.

The DRP Approach

Reorder point, and similar simplistic calculations, were moved out of the factories during the 1960s and 1970s by taking advantage of the dependent demand relationships expressed by the bill of materials. Material Requirements Planning (MRP) software tools were created to perform time-phased dependent demand propagation.

During the 1970s and 1980s a similar mechanism was proposed to move reorder-point out of the distribution channels. The distribution paths could be formulated into "bills of distribution." The same MRP logic could be applied to determine time-phased requirements of the parts in those distribution paths. The resulting re-adaptation of MRP logic is called DRP (Distribution Resource Planning).

Where it has been implemented, DRP has typically generated tremendous competitive advantage. However, it has not been widely implemented.

⁹ Andre Martin, *Distribution Resource Planning*, Revised Ed., Oliver Wight Pub., VT, 1993, p. 40.

DRP's Position in the Planning Process

The focus of DRP is on propagating demand from the customers, through the distribution network, to the factories. The demand at the factories is an output of the DRP effort. It is passed as input to the master planning efforts in those factories (assuming the traditional planning process, see Figure 2. - 4). When the factories have completed their planning cycles, their master production schedule becomes input to the DRP process. The schedule indicates when the orders will be satisfied. DRP propagates the supply flow from the factories back out through the distribution networks to the customers.

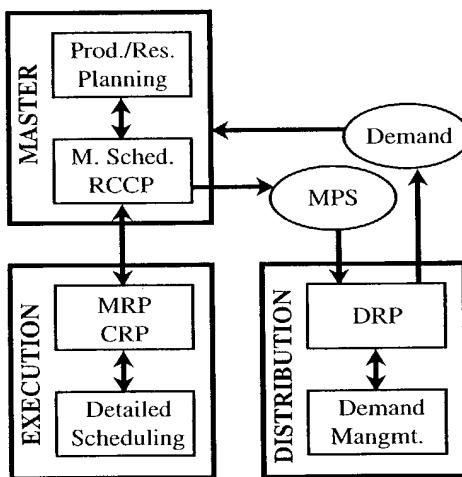


Figure 2. -4: DRP's Position

Note that when making distribution planning decisions, the planners do not see the supply. They can only guess what a factory can produce. When making production planning decisions, the planners do not see the distribution network. The network is where much of the inventory and all the true demand lies. They can only guess what is needed.

Resolving the problems created by poor guesses requires iterations between master planning and distribution planning. In addition, with the traditional planning process, making master planning decisions involves iteration between material and capacity tools, and between master level and execution level planning tools. A combination of any number of iterations between the various tools may be necessary before a good and feasible plan is devised. However, zero or one feedback iteration is likely to be the norm. Therefore, the quality of the plan hinges upon the ability of the various planners to guess well.

Compatibility of MRP and DRP

MRP and DRP are fundamentally the same algorithm. The primary difference is that the "Bill of Materials" represents which items are used to make other items, whereas the "Bill of Distribution" represents which items at different locations are used to feed other locations. However, since the

algorithms are the same, the bills could be merged into one. One algorithm could perform the time-phased propagation of dependent demand from customers, through the distribution network, through the factory and the component items, back through the supply chain to raw materials. Technically, it is a straight-forward generalization of the backward propagation algorithms of MRP and DRP.

Despite the technical feasibility of such integration, members of the supply chain are rarely connected at this level. Commonly they are connected by demand orders (purchase orders, requisitions, etc.) at the master planning level. Although the DRP could be directly integrated with the factory MRP, and supplier factories' MRPs could be directly integrated with consuming factory's MRP, the master planning step of the process is usually placed between the two (see Figure 2. -5). The master planning efforts are effectively detached by static demand orders. The result is limited visibility and the need for significant manual iteration cycles to develop a good and feasible plan.

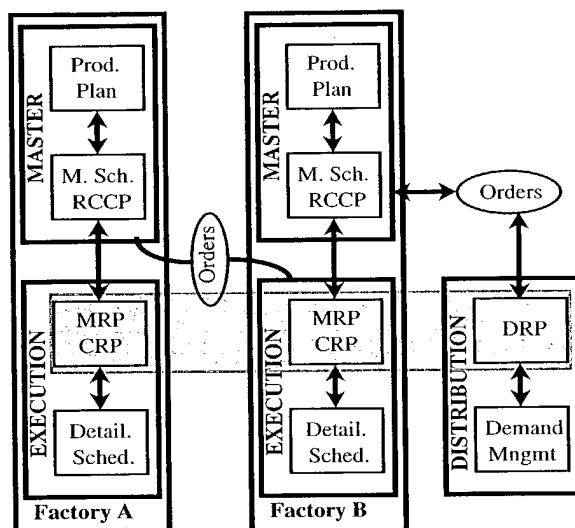


Figure 2. -5: United MRP and DRP

Limited Usage of DRP

Before proposing a TIP alternative to DRP, it is important to understand the current low penetration of DRP. Despite the acceptance of DRP as a superior mechanism, and despite that MRP has eliminated most reorder-point within factories, reorder-point and derivatives remain the dominant mechanisms in use in distribution networks. Why?

If a single member of the supply chain chooses to implement DRP alone, the gains are very small. They must forecast downstream demand and forecast upstream supply. The net result is only marginally better than reorder-point. The trouble may not be worth the gain. Improvement effort is typically poured into more advanced forecasting methods and more advanced reorder point variations.

Many members of a distribution network are often organizationally disjoint. No one person can make the decision to implement DRP for the network as a whole. Similarly, there is not a specific person or group that is held accountable for the performance of that distribution network. There is no direct incentive for someone to repair it. In contrast, a factory typically has a person who can unilaterally decide to implement MRP. That person is held accountable for the performance of that factory.

Furthermore, the distribution members are typically remote from each other. DRP software tools must be more sophisticated (electronic communications and related issues), as well as more flexible (to accommodate various organizational needs and preferences) than MRP software.

The Modern World

New manufacturing practices and new competitive pressures are having significant impact upon the relationships of the members of supply chains. This impact must be accounted for, and supported, by the planning processes that address the supply chain.

The Impact of JIT

The principles of Just-In-Time (JIT) include the importance of close relationships with suppliers. JIT calls for ordering just the items needed when they are needed. That typically means more and smaller orders and much more critical (unbuffered) due dates. Handling more and smaller orders is made easier by other JIT concepts such as elimination of purchase orders, incoming inspection, and other per-order overhead. For all this to work, the supplier must be willing and able to work in such an environment.

One result is long-term contracts that effectively authorize many purchases distributed over a time horizon. Such "blanket orders" must be supported intelligently by the planning system. The many individual due dates must be planned independently. But in doing so, the planning systems should not enforce excess paperwork or overhead to deal with the multi-order contracts.

More fundamentally, the planning system must be able to develop detailed plans that can be followed reliably (feasible plans). Vollmann, Berry, and Whybark note, "JIT's primary emphasis is on execution activities. Good planning precedes excellence in execution."¹⁰ Intelligent planning is an important ingredient in implementation of JIT.

Time-Based Competition

Time and responsiveness are the competitive drivers of this decade. JIT and quality are prerequisites. The ability to react quickly to changes in demand is critical. As the marketplace becomes global, manufacturers must provide high-levels of options in easily customized items. Short lead times are also very important. To support this, inventory throughout the supply chain must be minimal. The material must move through the supply chain rapidly.

¹⁰ Vollmann, Berry, Whybark, *Manufacturing Planning and Control Systems*, Third Ed., Business One Irwin, 1992, p. 101.

The elimination of inventory and acceleration of material movement (smaller queue times) effectively removes all the buffering that once protected systems from poor planning. If a manufacturer assumes large queue times and safety stock, then small plan infeasibilities are typically not problematic. The slack in the system can be consumed to make up for weak planning. Without that slack, planning mistakes become critical. Time-competitive, low-inventory systems cannot be maintained without accurate and intelligent planning.

Modeling the TIP Supply Chain

TIP requires immediate visibility of all important effects throughout the supply chain of any planning decisions. That does not mean that models of the different members of the supply chain need to be in the same level of detail. In fact, organizational constraints may prohibit the same level of detail. Furthermore, planners in the various organizations are planning for different reasons, and may need different information.

Visibility

If the whole supply chain in Figure 2. -2 was a single JIT company, the excess inventories would be unacceptable. As a collection of JIT companies, those excess inventories should be just as unacceptable. In both cases, the competitiveness of the supply chain suffers. And in both cases, the same visibility and intelligent planning is required to have any real chance of eliminating it.

Ideally, providing such visibility is relatively straight-forward: expand the factory model to include the entire supply chain and regularly obtain status data from the other members of the supply chain. Reality, of course, is not so simple.

Point-of-Control

Each member of the supply chain may provide visibility to the other members, but they may not necessarily want to provide authority to change their plans. Different supply chains need to be controlled differently. Some are controlled by a central planning organization. For other supply chains, control is distributed throughout. The majority are a mix of the two. The issue of authority must be supported by TIP software.

Centralized Control

Given the global visibility of the supply chain provided by TIP, it is possible to have global control over the supply chain. A planner could see the forecasts and customer orders from all the sales organizations. The inventories in the distribution centers and warehouses would be visible, allowing coordination of the supply of those orders through that distribution network. The capacity and status of the factories would be visible. As a result, the demand could be distributed evenly on those factories avoiding overloading or underloading any one of them.

When possible, such centralized control over the supply chain typically provides the best results. TIP software must be able to support this type of

control over the plans of multiple members of the supply chain wherever possible. Where it is not possible, however, TIP software has a more difficult challenge: *distributed control*.

Distributed Control

The members of the supply chain are typically different companies with different goals. Few companies are going to agree to submit their business planning to some external central authority.

Although the need for distributed control is more obvious with separate companies, it is also often necessary for a single company. Large companies are typically made up of separate organizations or business units, which have their own goals and are evaluated against different criteria.

People cannot be held responsible for their organization's performance if their operations are planned by others. Furthermore, the expertise for planning different operations almost always exists within the individual organizations (not in some corporate planning organization).

Although TIP calls for *visibility* of the whole problem, it in no way calls for *global control* over the whole problem. There must be explicit support for separate "points-of-control" in the global model. A TIP system must support having portions of the model that a planner *manages* and portions of the model for which only visibility is provided. A planner must be able to see the effects his decisions have on the entire supply chain, but he can only make planning decisions in the portion of the chain over which he has control.

Hierarchical Control

When multiple members of the supply chain are part of a single company, some level of centralized control is possible and desirable. However, it is typically not complete control.

For example, the central planners may be coordinating distribution of supply and demand through the supply chain, but they are not planning the individual members. The sales professionals may be in control of the forecasts and customer orders for their market area. The factory planners may have full control over their factories. However, the selection of which factory supplies what demand from the market areas may be coordinated by a centralized planning organization. Using their visibility of the capacity and status of factories, they can intelligently and cooperatively route demand to the factories.

Often, many alternative flows in a supply chain can fulfill an order. Choosing a flow path may be one of the most complex planning decisions made. The effectiveness of a particular choice cannot be known until a plan has been developed for it. Evaluating alternate flows involves developing alternate plans for each and evaluating the differences.

Because it may cost more to supply from distant or alternate factories, evaluating alternative flows is particularly important for large or multi-national companies with distributed capacity. Profitable demand should not go unsatisfied because the planning process only allows one factory to serve a market area. Different factories may be more adept at certain products, but

if demand for those products is low, they should not be artificially restricted from building products that are experiencing higher demand.

In some large companies, distributing demand among multiple factories is as fundamental a planning decision as distributing load among resources within a factory. To do that effectively, the corporate planners must see all the plants and their demands. Although the corporate planners are distributing load, the individual factories' planners must retain authority over their respective factories. The corporate planners cannot know sufficient detail to plan the factories themselves. If the factory personnel are to be held accountable for the factory's performance, they must have control over the plan that is driving it.

The definition of the "points-of-control" must be flexible so that centralized control can be employed wherever it is beneficial. It must not be limited by geography, function, or other artificial division.

Public vs. Private Information

Each member of the supply chain must supply the model of its internal operations and the current state of activity and plans to all the other members of the supply chain. If the supply chain was one company, then there would be little problem sending all the data concerning the supply chain to the individual members. When it is not one company and different members may be competitors, simply transmitting all operational data is completely unreasonable.

However, supporting TIP does not require detailed operational information about each of the companies. Each organization simply needs to be able to model how their decisions affect the other organizations. An organization only needs to publish a detailed enough model to allow others to properly model their effects. The more detail they publish, the more accurate others' visibility is. Each organization can independently decide what to publish and what to keep private (see Figure 2. -6).

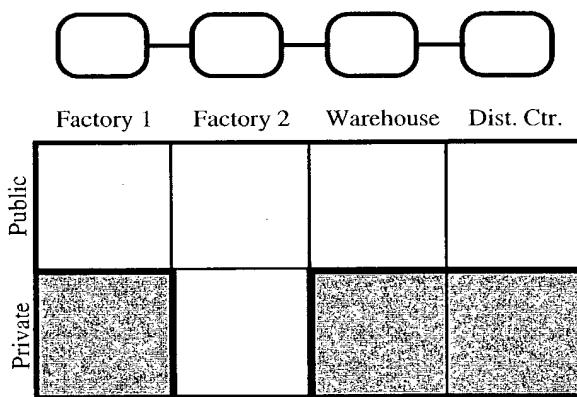


Figure 2. -6: Supply Chain Privacy

Separation of public and private data in the models must be supported directly by TIP software. There should be only one model of each

organization. Maintaining more than one model always leads to inconsistencies and headaches. That one model must be useful internally as the detailed model, and in filtered form as the published model.

Various Levels of Detail

A corollary of the previous discussion is that different portions of the model of the supply chain are in various levels of detail. Your own organization needs to be modeled in enough detail that detailed planning decisions can be made. Sister plants and distribution organizations may also be modeled in fairly detailed form to facilitate high levels of cooperation. Organizations in other companies may provide much less detail (or you may want less detail). This last point should be emphasized: even though an organization provides a great amount of detail, the other organizations are not obligated to use that detail. They are each free to further simplify the published models of the other members (see Figure 2. -7).

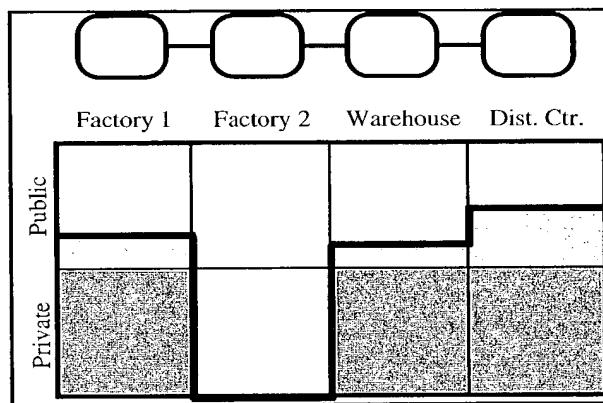


Figure 2. -7: Model Simplification

Even with the tremendous computer power available today, modeling the entire supply chain can be computationally expensive. TIP calls for immediate visibility of the effects of planning decisions. However, immediate visibility cannot be provided if the model is too large and complex. Eliminating the detail that is unneeded for a particular planner's decision making is important for complexity reasons as well as computational reasons.

Many software packages are not designed to support multiple levels of detail. The software is designed for some particular level of detail, and less detail is provided by filling in zeroes (or default values). Such software technology is inadequate for TIP. The software must be able to efficiently take advantage of reduced detail to provide immediate visibility.

Pruning the Network

Not only must the level of detail in external organizations' models be reduced, but also much of the global network should be pruned away. The network that a planner needs is the collection of supply chains for all the

products fed by that planner's operations. For example, the fact that ABC Manufacturing is fed by XYZ Producer (see Figure 2. -8) does not mean that every supplier and every customer of ABC need to be modeled by XYZ. The effect that planning XYZ has on such organizations may be too indirect to be of importance in making those planning decisions. TIP does not call for each manufacturer to model the world. Rather, TIP calls for each planner to see any important effects that their planning decisions have on the supply chains of interest.

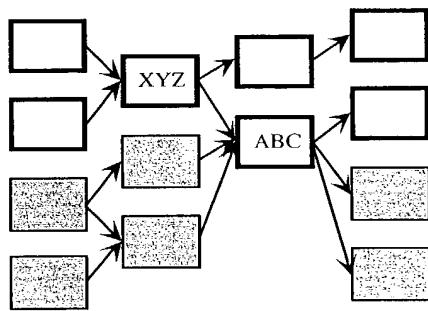


Figure 2. -8: Pruning

Connecting the Network

The net effect of modeling a supply chain is as if there was one combined bill of distribution and bill of material (BOM) describing how an item is built and distributed. In reality each organization has its own portions of that virtual BOM, each with their own different part numbers. Supporting connection of models in the supply chain requires transparent support for mapping between different organization's models and naming schemes. Planners must have easy, clear access to that mapping. Communication between the planners in different organizations relies on clear access.

Global Visibility and Local Focus

Besides the separate points-of-control and various levels of detail, different organizations in the supply chain are fundamentally different in their operations. There may be make-to-order, assemble-to-order, and make-to-stock manufacturers in a single supply chain. There may be continuous, repetitive, and batch manufacturers in a single supply chain. There may be miners, manufacturers, warehouses, transportation companies, wholesalers, retailers, and sales organizations in a single supply chain. In addition, there may be many different industries (e.g., automotive, steel, textile, chemical) represented in a single supply chain.

Historically, software vendors have provided different tools for companies in different industries, for different types of manufacturers, and for different types of organizations. A software package intended for batch manufacturers may have no support for rate-based manufacturing. Make-to-stock packages may have no order configuration tools.

TIP software packages must be able to represent all these type and focus differences within one model. Rate-based models must be intermixable with

discrete models. Models of chemical, metal, paper, and automotive industries must be able to merge into one model. Each organization must be allowed to choose its own focus and orientation, and it must be well supported by the tools. The differences in focus are real and important. TIP adopters should not have to sacrifice local focus to gain global visibility.

Application of TIP to the supply chain becomes infeasible without software that can model different industry segments, different manufacturing types, and different organization types. Furthermore, many individual manufacturers exhibit traits of many different manufacturer types and industries. Truly integrated planning of such organizations, as with supply chains, requires software that can model it all.

Driving the TIP Supply Chain

The distributed nature of the supply chain affects not only the model but also the process itself. Many decisions about the supply chain cannot be made locally. There must be a way to plan locally using visibility of others but without forcing decisions on others.

Propagating Supply and Demand

An effective mechanism for propagation of supply and demand information is the key to providing full visibility of the supply chain. The traditional process propagates demand by sending purchase orders to the supplier. Supply is similarly propagated by the acceptance of the order and assignment of a due date by the supplier. However, that is more of a reactive mechanism than a planning mechanism. Since it affords no visibility to either supplier or consumer of the effects of making changes to those orders, it is a relatively static mechanism.

If, instead of a simple purchase order, each consumer sent to the supplier a model of how and when the orders were going to be used, then the supplier could evaluate the effect of missing a due date. For example, one of the supplier's machines unexpectedly goes down for one day. Which orders should be delayed? Which orders can be delayed with least impact on the final customers? Perhaps one of the consumer's machines has gone down, and the orders are not even needed at the original due dates. These questions are not answerable with simple purchase orders.

Providing such visibility to planners is the foundation of Truly Integrated Planning. A planner should have a planning model of all significant upstream suppliers and downstream consumers. When making a planning decision, the effects on those upstream and downstream members should be immediately visible.

Note that the object is not to plan those other organizations. Rather, it is to provide visibility of the effects of local planning decisions on the supply chain as a whole. That may involve emulating the typical planning decisions that are made at remote organizations in response to the decisions made locally. But only rough estimates are needed from these organizations. For the majority of the situations, much better than rough estimates can be provided (particularly if the organizations are using consistent software).

tools). Global visibility provides the foundation for intelligent planning decisions.

The Customer Interface

Although not always considered as such, the act of accepting an order and assigning it a due date is a planning activity. In fact, it can be argued that setting order due dates is the single most critical planning activity. Of all the dates set in the planning effort, the customer due dates are the only ones that the company has a commitment to fulfill. Gopal and Cypress observe, "Order management systems provide the front-end to the integrated supply chain. As such, they enable the primary customer interface and drive the entire supply chain."¹¹

The majority of independent points-of-control in many supply chains are the sales and order entry points. These points, as interfaces to the customers, need accurate up-to-date visibility of the manufacturing and distribution networks upon which they are completely dependent. Those salespeople must be responsible for the negotiations and commitments given to their customers. In such a capacity, salespeople have a significant role in the planning of a supply chain's operations.

The planning effects of sales personnel are typically managed by the planning personnel through mechanisms such as available-to-promise (ATP). However, ATP is not a very dynamic mechanism. In contrast, sales is typically a very dynamic environment. TIP provides the opportunity for far more sophisticated order quoting mechanisms. Such mechanisms are more dynamic and allow the full capacity of the supply chain to be used by the sales people in meeting their customers' requests.

By providing such visibility and order promising mechanisms, the TIP process facilitates much higher responsiveness to customer demand. As the real demand arrives, divergent from forecast and business plans, planning decisions can be immediately initiated to reallocate capacity to follow the demand pattern. That is the topic of the following section, "Integrating the Management of Supply and Demand."

¹¹ Gopal and Cypress, *Integrated Distribution Management*, Business One Irwin, 1993, p. 160.

3. Integrating the Management of Supply and Demand

This section discusses the reengineering of Demand Management, fully integrating it with Supply Management to form a single Truly Integrated Planning process. Without such integration, global visibility of the planning problem is not available, and global supply chain optimization is not possible.

The preceding sections “Reengineering the Manufacturing Planning Process” and “Integrating the Expanded Manufacturing Planning Process” introduce a planning process reengineered to take advantage of modern assumptions about technology and the manufacturing world in which we live. Those process changes have significant impact on demand management, and this section discusses that impact.

This discussion focuses on the reengineering of Demand Management, fully integrating it with Supply Management to form a single Truly Integrated Planning process. Without such integration, global visibility of the planning problem is not available, and global supply chain optimization is not possible.

Promising is Planning

Although not always considered such, the act of accepting an order and assigning it a due date is a planning activity. It can be argued that setting order due dates is the single most critical planning activity. Of all the dates set in the planning effort, the promised due dates are the only ones that the company has a commitment to fulfill. Gopal and Cypress observe¹², “Order management systems provide the front-end to the integrated supply chain. As such, they enable the primary customer interface and drive the entire supply chain.”

Thus, in some sense, the majority of independent planners in many supply chains are the sales and order entry points. These points, as interfaces to the customers, need accurate, up-to-date visibility of the manufacturing and distribution networks upon which they depend. The salespeople are responsible for the negotiations with and commitments given to their customers. In such a capacity, salespeople have a significant role in the planning of a supply chain’s operations.

¹² Gopal and Cypress, *Integrated Distribution Management*, Business One Irwin, 1993, p. 160.

The planning effects of sales personnel are typically managed by the planning personnel through mechanisms such as available-to-promise. (Available-to-promise is discussed later in this section and in the section "Reengineering Order Promising to Satisfy Business Objectives.") However, traditional ATP and traditional demand management are not very dynamic. In contrast, sales is typically a very dynamic environment.

Truly Integrated Planning

Based on a significantly different set of assumptions, Truly Integrated Planning (TIP) involves many fundamental changes from the traditional planning process. The tools and process steps are *integrated* into one. The planning problem is addressed as an *integrated* whole, allowing global optimization. Material and capacity issues are *integrated*. Detailed near-term issues and longer-term issues are *integrated*. Factory and distribution issues are *integrated*. An *integrated* picture of the problem provides instant and global visibility of the effects of planning decisions throughout the supply chain. The on-going planning efforts are *integrated* over time to effectively evolve better plans. Automated and manual planning are cooperatively *integrated* so that they can each be used to their maximum potential. That is the level of *integration* called for by *Truly Integrated Planning*. Separate tools that have simply been interfaced together do not satisfy TIP requirements.

Traditional Dis-integration

A traditional order entry system allows the order to be recorded. The order is then fed through data interfaces or a common database to the various planning systems. A few days (or longer) later, a plan for that order is complete. Hopefully, the resulting plan satisfies the order.

A "sophisticated" ATP order entry system might allow the order to be immediately given a due date. However, if the ATP does not satisfy the order, the order is either rejected or goes through the dis-integrated process above.

Visibility of Sourcing Options

With a Truly Integrated Planning system, planning a new order can instantly disseminate the change through all aspects of the plan. When a new order is planned, materials are consumed, both factory and distribution capacity is reserved, near- and long-term effects are identified, and a full, feasible plan may be developed.

Thus, evaluating sourcing alternatives for satisfying a new order is straightforward. The distribution routings identify the alternative sources, allowing sourcing from any of the plants that can produce the item. In each plant, the manufacturing routings identify the alternative methods to produce the item and the alternative components or suppliers that can be used.

Visibility of Cost Implications

After the feasibility of the plan, one of the primary issues in evaluating all the sourcing options is the cost. For example, if a new order with a profit margin of \$500 is being considered, it may not make sense to satisfy that order by selecting a sourcing option that costs an extra \$800, unless the price to the customer is raised comparably.

Traditional planning and order quoting systems do not have visibility of such options and costs. However, without such visibility, how can the decision to accept an order be made intelligently? How can an order be quoted with a certain price if the cost to satisfy that order is unknown?

Traditionally we have resorted to limiting sourcing options and averaging the costs. Such hardly qualifies as supply chain optimization. In today's world where responsiveness is key, such rigid mechanisms quickly become a competitive ball-and-chain.

Visibility of Allocation Options

For any sourcing option with resources in scarce supply, TIP requires immediate visibility of the effects of allocation decisions. For example, if one of the operations planned to use a resource must be delayed, what demand is affected? What is the cost of not satisfying that demand?

Evaluation of sourcing options often leads to evaluation of allocation options. If a high-priority request comes in, satisfying that request may mean canceling some forecast orders (future sales). Is this sale higher profit than those future sales? Without knowing the answer, an intelligent decision of whether to promise that request cannot be made.

Critical Optimizations May End with a Promise

The decisions that may have the most impact on supply chain optimization are order promising decisions. In fact, most of the cost versus profit margin trade-offs must be made prior to promising. Once the promise is made, the cost issues change!

Before an order has been promised, it makes sense to say, "I'd have to run overtime to satisfy this request, and that would cost more than the profit I would make, so I'll reject this request."

In contrast, if the order has been promised, the cost trade-off is different. There is a commitment to satisfy the promise, and failing to satisfy it may have costs far greater than the lost profit margin. That same order above that would have been rejected *prior* to promising, may result in overtime if the need is not identified until *after* promising.

TIP Must Include Demand

Planning tools are traditionally designed to develop a supply plan to satisfy the demand inputs. However, the real planning problem (for which the planners should be held accountable) includes planning *both* the supply *and* the demand. Options are exercised on both sides, balancing supply and demand in such a way that the business objectives are achieved. The biggest

opportunity for cost optimization, *prior* to order promising, must not be ignored.

Thus, Truly Integrated Planning must include the demand management activity as well as supply management if true global visibility is to be provided. Global supply chain optimization is not possible unless both supply and demand options are visible and exercisable.

Traditional Demand Management

As stated earlier, the planning effects of sales personnel are typically managed by the planning personnel through mechanisms such as available-to-promise (ATP). However, traditional ATP is not very dynamic, while sales is typically a very dynamic environment.

ATP for Quick Promising

In environments where customers are not willing to wait for a plan to be developed to get a promise, the supplier must plan in advance. The supplier must create reliable promises that are available for immediate transfer to a customer. The future requests can be forecasted and a plan can be created to satisfy and promise those forecasted requests. Those promises to forecasted requests are “available to promise” to customer requests.

When an actual customer request is received, one or more (or a portion of) promises made to forecast requests may be instantly reassigned to the customer request. The forecast requests are “consumed” by the actual requests.

Even in environments where absolute speed is not essential, high-volume promising may necessitate an ATP mechanism. Consider an organization with 100 people taking orders simultaneously. Having each of the 100 simultaneously manipulating the same plan, competing for available capacity and materials, is a cumbersome solution. It is much more efficient and structured to compute ATP for the various products, and then let the salespeople simply compete for that ATP. (Or they may compete for allocated portions of that ATP, as is discussed later.)

ATP for Proper Promising

The sales people preparing the quotes are usually not the factory planners. The sales people should not be making planning decisions. And yet they may be setting the most important due dates of all. This dichotomy has traditionally been managed via ATP.

The planning organization is responsible for taking in forecasts from the sales organizations and then producing a master production schedule and ATP. That ATP is then handed to the sales organization for making promises to customers. The promises made by the sales people are limited by the ATP given to them by the planners. The sales people can only request more ATP.

ATP Reduces Flexibility

The disadvantage of the traditional ATP process is its static nature. ATP numbers are generated as part of the master production scheduling effort performed monthly, or sometimes weekly. Each planning cycle begins by revising the sales forecasts. Those are passed to the planners who prepare a new master schedule that is based on those revised forecasts and the current backlog of actual requests. That master schedule minus the backlog is the ATP that is available for the rest of the planning cycle.

Such practice often is simply not responsive enough to keep up with our ever-changing world. Sales is too dynamic an endeavor to be limited by such inflexible numbers. If sales do not follow forecast, potential sales can easily be lost and capacity can be left idle unnecessarily. Given the difficulty of forecasting (guessing the future), all performance dependencies upon forecasts should be minimized.

Forcing stability by sticking with the monthly ATP numbers can lighten the planners' load, but it leads to turning away customers ready to buy while reserving capacity for sales that may never materialize.

Dynamic Demand Management

TIP must support sophisticated order quoting mechanisms that are more dynamic. The full capacity of the supply chain should be available to the sales people for meeting their customers' requests.

By providing such visibility and order promising mechanisms, the TIP process makes it much easier to be highly responsive to customer demand. As the real demand diverges from forecasts and business plans, planning decisions can be immediately initiated that reallocate capacity to follow the actual demand pattern.

Making ATP Dynamic

A fundamental principle of TIP is that planning is *not* the process of building a plan. It is the on-going process of maintaining a plan. ATP becomes dynamic when it is placed in a dynamic planning process.

All customer activity should be visible. (That may mean modeling downstream organizations in the supply chain.) TIP couples visibility with an understanding of the expected demand pattern. Rules can then be set up to automatically adjust the forecasted requests in response to the incoming actual requests.

ATP for a monthly period is generally not consumed suddenly. As time passes, more and more of it is gradually consumed by actual sales.

The planners, working in an on-going process, see the sales activity as it gradually consumes the ATP. They can see in advance where the ATP is likely to run low, and where it is more than adequate. They can re-evaluate forecasts as actual orders do not follow them.

Given such visibility, the planners generally have plenty of time to adjust the build plans to shift ATP to where it is needed. They can do this with full

consideration of which products are more profitable, of which sourcing options are more expensive, and so on.

Effectively, the planners can make the critical planning decisions about which orders to promise and which to reject in advance. Those decisions are recorded in the form of ATP. Ideally, the ATP is always adjusted properly before it is actually needed by the sales people.

When ATP is Not Enough

Even with an effective and dynamic planning process that has full visibility of both demand and supply management issues, customers still occasionally surprise the planners. Thus, there are some cases where the available ATP does not satisfy the customer request.

With an on-going TIP process, the sales person informs the planner that a special order has been received that needs immediate attention. The planner has full visibility of both supply-side sourcing options and of demand-side re-allocation options. For example, to satisfy a special order for widgets, some of the remaining ATP for gadgets may be reduced, thereby freeing the necessary capacity to build the widgets.

Given a process that maintains dynamic ATP, support for rapid turn-around on special orders becomes feasible.

Dynamic Forecast Management

Although the fine art of forecasting (predicting the future from the past) may be performed once each month (cycle), forecast management should continue throughout the month. For example, if 3000 units is forecasted for the next month, and the demand is expected to arrive fairly evenly throughout the month, then forecast management can adjust the forecast if reality does not follow. Roughly 100 units per day are expected. If after 10 days have passed only 500 units have been sold, it may make sense to adjust the forecast from 3000 down to 2500 (assuming the rest of the month will still sell at 100 units per day) or down to 1500 (assuming 50 units per day will continue).

Similarly, if 2000 have been sold after the first 10 days, and demand through the month is expected to be fairly level, then it may make sense to adjust the forecast from 3000 up to 4000 (assuming the rest of the month will still sell at the original forecast of 100 units per day) or up to 6000 (assuming 200 units per day will continue).

Note that such forecast adjustments need not involve redoing the "forecasting" computations (re-examining the historical numbers). Rather, it simply involves a rule for how to adjust the forecast and the numbers computed previously by the forecasting effort. The term "forecast management" is used to describe these dynamic adjustments to forecasts outside the traditional "forecasting" effort.

Dynamic forecast management rules can specify how forecasts expire. The forecasts may be shifted out in time if they are not consumed; others may expire and disappear if not consumed. Such auto-adjustment of forecast numbers can be very valuable in maintaining accurate forecasts and allocations for hundreds or thousands of products between "forecasting" cycles.

4. Reengineering Order Promising to Satisfy Business Objectives

This section describes a reengineered Order Promising process based on "Allocated ATP". That is ATP that has been allocated to particular sales people or organizations, to particular customers or market segments, and/or to particular products or product characteristics. The result can allow very flexible on-the-phone order promising that is backed by a reliable plan and that results in satisfying detailed business objectives.

How quickly can you get what customers want, the way they want it, and when they want it? That is the competitive question of the 90s, and an answer is often expected while still on the phone. And a reliable answer is expected, requiring there to be a solid plan to back it up.

The traditional answer of "I'll get back to you with a quote in 3 days" could be 3 days after your competitor has signed the deal. Reliable answers can be pre-planned in the form of ATP, but traditionally, flexibility is lost. Sales that could have been made are turned away, while capacity remains reserved for things that are unwanted. Quick, unreliable answers are an easy way out, but such gambles rarely pay off in the long run.

Dynamic ATP is needed to be responsive to customer demand. Achieving this was discussed in the section, "Integrating the Management of Supply and Demand." Taken to an extreme, quality planning tools can be set up such that the ATP automatically tracks customer demand. The ultimate in "customer-driven planning" is when whatever customers are buying most, the manufacturer builds the most. However, although it is a popular phrase, do we really want to be "customer-driven"?

Who's Driving?

Although planning tools should be *capable* of supporting a completely customer-driven process, for most businesses the process should *not* become completely customer-driven. The major planning decisions should not be made by the sales process or the customers. Order entry personnel should not be changing the plans for the factories in the supply chains, nor should they be changing the business plans for the organizations involved.

Given an integrated planning system with a global view of the supply chain, its order promising system can take a new order, map it into available capacity, and quote the earliest possible delivery date. The due date quoted is

reliable (since it is based on a feasible schedule) and as competitive as possible.

Such a process *could* provide the ultimate in “customer-driven” planning. However, the resultant plans would not always be desirable. The effect would be that a business’ capacity would be allocated on a first-come, first-served basis, with no regard to the other business factors that need to be evaluated.

For instance, consider a business that focuses on making “widgets”, a high-margin product with volatile demand patterns. In addition, the business also manufactures “gadgets”, a low margin, but high-volume commodity product. When demand for widgets is low, the high demand for gadgets is allowed to consume the excess capacity (better to make low profits than to lose money).

Given the “customer-driven” planning process just described, the sales people sell to any customer as long as there is capacity. The high demand for gadgets quickly consumes most of the business’ capacity, well in advance. The sporadic demand for widgets has little opportunity to be filled. The net result is a significant loss in profitability.

Thus, despite the prevalence of “customer-driven” buzz words, it is important to recognize that much more than customer demand goes into the planning of what to sell. The planning process must not allow itself to become solely driven by that customer demand. The planning process must carefully balance these two facts: order promising is an important part of the planning process, but order promising cannot dictate the plans.

Traditional ATP-by-Item

The traditional method for controlling the planning effects of sales personnel is available-to-promise (ATP) logic. Sales organizations are given ATP numbers for each item, and they sell and promise due dates as dictated by those ATP numbers. The capacity of the organization has been pre-allocated to certain items. The sales organization must sell based upon that allocation.

ATP-by-Item solves the former problem: the widgets and gadgets are pre-allocated a certain level each period. Once the demand for gadgets has consumed the allocation in a period, further gadget sales are pushed off to the next period. That guarantees that the capacity to produce widgets, up to the quantity allocated in that period, is not consumed by lower-margin gadget sales.

Product-Allocated ATP

ATP-by-Item is a great start, but it is not enough. There are many other characteristics of the products that companies sell that can dictate margin and thus have need of protection via pre-allocated ATP. Some examples are product grade, delivery lead time, quantity, price, customer, and market segment.

ATP-by-Grade

Some products have a "grade" such that a higher-grade product can be sold as a lower-grade product. If the demand for the higher-grade product is less than is produced, then it can make economic sense to sell the higher-grade product as if it was lower-grade (this is often called "downgrading" or "downbinning").

Consider that widgets are divided into three grades depending upon their "speed", as determined by post-process testing. The "speed" of widgets is not entirely controllable, and thus the yield mix of different grades from the manufacturing process is highly unlikely to match the product mix desired by sales. If you get an order for fast widgets, you cannot fill it with slow widgets. But if a customer requests slow widgets, they will not complain (or even notice) if they are given fast widgets. Thus, downgrading of fast widgets is possible.

However, care must be taken. It may be simply that the demand for the higher-grade widgets will come later in the period. Given that the higher-grade widget is also higher-margin, it makes sense to downgrade only if the sales of the higher-margin widgets will not consume all that is available.

Thus, fast widgets have the same problem that widgets do: if left solely to customer-driven mechanisms, the high-volume demand for slow widgets completely consumes all of the fast widgets, resulting in lower profit margins. To prevent this, ATP for widgets can be split by product grade. That prevents the allocated level of fast widgets from being consumed to fill orders for slow widgets. This can be accomplished by giving the different grades separate item numbers (which may be wise for other reasons as well). But the ability to downgrade must not be lost.

ATP-by-Delivery-Lead-Time

Customers are often willing to pay more when they need something immediately. For example, the normal price for a gadget may be \$10, with a delivery lead time of 6 weeks. However, a customer may be willing to pay \$20 if you can deliver in 2 weeks; or \$30 if you can deliver tomorrow.

The profit margin differences can be huge, especially if competitors cannot offer similarly short lead times. Of course, the cost to produce with such short lead times may be higher. (Typically it means that more inventory must be kept later in the process.)

Since 6-week lead time demand always arrives several weeks prior to 2-week lead time demand, it gets first shot at the ATP for that item. However, the manufacturer prefers that the 2-week lead time demand have precedence since it is higher-margin.

Effectively, the salesperson must gamble -- how much short-lead-time demand will there be? If the sellers guess high, they end up selling less product than was possible; if the sellers guess low, they lose high-margin sales to lower-margin sales.

Once again, pre-allocated ATP can be the answer. By separating out the ATP for short lead time products, the lower-margin sales are not allowed to consume it.

ATP-by-Quantity

Another common product price split is based on quantity discounts. Consider an industry-standard price break at 100 gadgets, where about half of the normal gadget sales are more than 100 gadgets. The price break at 100 is necessary to be competitive and not lose sales, but it is undesirable to sell large quantities in lieu of smaller, higher-margin orders.

Once again, the sales of smaller orders can be protected from the lower-margin orders by separating out the ATP by the quantity discount.

ATP-by-Price

The pattern in these examples has become obvious. The particular product characteristic used to determine margin is not the issue. Good ATP software should allow ATP to be separated based on *any* criteria that affects profit margin. In fact, the criteria can be as simple as the price. It is not unheard of for a supplier to provide a certain amount of discount product, first-come-first-served or otherwise.

Splitting out all these margin-oriented products using separate item numbers can be cumbersome. First, it forces the customer to choose the appropriate item (discussed more later in this section). Second, it artificially reduces flexibility. Given the natural fluctuation of demand mix, it is critical that ATP be applicable to as many different orders as possible, and that ATP can smoothly transition to less restrictive products.

Market-Allocated ATP

So far all ATP allocations were made to protect higher-margin products from being consumed by lower-margin ones. There are many other valid reasons to allocate ATP. Common splits are by customer and by market segment.

ATP-by-Customer

Gadget International may be a 30% shareholder and thus get preferential treatment. In fact, they may be guaranteed a certain number of gadgets each month, though they may only have to buy the amount they want.

In such a scenario, it is important to pre-allocate enough ATP to Gadget International to guarantee that other sales do not consume the capacity that is guaranteed to them.

On the other hand, there is probably a due date for requests from Gadget International. After that date, it is important that the ATP that was exclusive to them be freed for consumption by others. Splitting these two products into separate item numbers can be cumbersome.

Widget Pro Inc. may be a special VAR (value-added reseller) that gets different pricing and delivery lead times. It is necessary to split out the ATP for those products from the products that are sold to other customers (even though they are all the same item, widgets).

ATP-by-Market-Segment

Market segment or sales channel is another common sales split. There can be many reasons that sales to one segment or through one channel may be treated differently.

For example, Europe may be considered a more lucrative region for widget sales. To facilitate expansion into Europe, a certain level of sales are allocated to that segment, thereby preventing the currently larger US widget sales from consuming all capacity. In fact, to further stimulate expansion into Europe, pricing may be adjusted to be more competitive. European customers may be offered pricing that is not offered to US customers.

Similarly, different sales channels, such as retailers, mail order, phone order, direct sales, VARs, and third-party consultants, may be targeted for stimulation or special treatment.

By allocating ATP for a product to a particular market segment or sales channel, those sales can be protected.

Note that the alternative of assigning separate item numbers is undesirable. The demand from these markets varies, and it is important to be able to easily move ATP from one product bucket to another. The world of sales does not tolerate rigidity.

Consider if Gadget International requests more than its guaranteed amount. Surely refusing them blankly is not the right answer. Rather, they should have as much right to the remainder as any other customer (possibly even more right). If Gadget International requests 500 gadgets, they should not have to re-issue the order with a different item number when they exceed their guaranteed allocation. This flexibility issue is discussed more later in this section.

Seller-Allocated ATP

Most manufacturing companies have more than just one salesperson, often distributed across the country or globe. Thus, for most manufacturing companies, a single ATP number for each product is often inadequate.

A key purpose of ATP is to allow a salesperson to make a reliable promise quickly. If hundreds of people across the globe have access to the same ATP number, then they must take turns consuming from it (promising). If there are 1000 units left and there are 10 different salespeople who have received requests for 400 to 800 units, each salesperson thinks they can make a promise. But if they all do, less than a quarter of those promises can be satisfied!

Forcing all salespeople across the globe to take turns promising would probably not provide the necessary quick response. There are also the potential logistics problem of having them all connected to a single computer for each sale. One way to resolve this is to break up the one ATP number and allocate it to the various sellers. In that way, each owns a portion of the ATP and can promise from it independently.

ATP-by-Seller

There is another good reason to allocate ATP to individual sellers. ATP is based upon forecasts. Forecasting is the responsibility of salespeople. Tom Wallace argues¹³, "The best people to do the forecasting are those who will be held accountable for achieving the forecast."

By allocating the ATP to the sellers who forecasted the demand, those sellers are held accountable. Sellers that forecast too little are allocated less than they could have sold (usually meaning less commission). Sellers that forecast too much are not able to sell all that they were allocated and have the normal ramifications.

Of course, there is a flip side to this issue. If ATP is allocated to sellers, it should also be dynamically re-allocated. The nature of forecasting is such that sellers are normally incorrect—one sells less, another more. That is where sales organizations often come in.

ATP-by-Sales-Organization

Sales people are often organized into sales offices or organizations. At a single sales office, having all the sales people working on one ATP computer system is reasonable. Thus, it may be unnecessary to separate out the ATP.

This is particularly true in the case of order-entry personnel. The ATP can be allocated to the sales office, and then the order entry people that are members of that sales office can consume its allocation first-come-first-served.

On the other hand, the sales people may be responsible for their own forecasts. In that case, they may want their own ATP allocation. After all, if Jill accurately forecasts 100, it is not fair if she gets less because Jeff forecasted only 10, but actually sells 50.

Allocating the promises made to forecast requests to specific sales people allows them to use their own ATP numbers as they see fit. They can sell ATP without concern for whether they are using the same forecast promise as another seller, and without needing to check with each other before making promises.

Allocation Hierarchies

Even in the case where the sales people want their own ATP allocations, they may not want to be limited to that exclusively. To do so requires them to be perfect in their forecasting.

Forecasting is much more accurate in aggregate, and thus the forecasts for the sales office as a whole are generally much more accurate than each sales person's forecast.

A common approach is to have each sales person produce two forecast numbers: one that is what they think they can sell and a second that is what they are willing to commit to selling. ATP is then allocated to them based upon what they are willing to commit to selling.

¹³T. Wallace, "World-class Order Fulfillment: Part II -- Information Please", *APICS - The Performance Advantage*, November 1995, p. 64.

If the sales people do a good job they should all exceed their allocations, although which sales people will exceed by what amount each period is completely unknown. The sales office itself will want to commit to more than the sum of its sales people, since several will have likely committed to less than they can actually sell. Thus, the sales office gets a larger ATP allocation than it allocates to the individual sales people.

As the sales people consume all of their own ATP, they can begin to consume from the allocation that was retained by the sales office. This can be done first-come-first-served, or the sales office can increase the individual sellers' allocations as they near their limits.

Note that this same pattern can be taken up another level. Sales people are often organized into sales offices which are organized into sales regions which are organized into sales divisions or business units. Such organizational hierarchies lead to the need for ATP allocation hierarchies.

The sales offices' commitments can be handled by the sales division in much the same way that the sales people's commitments were handled by the sales office. As each sales office approaches its ATP allocation limit, it may go to the sales division to consume from the ATP allocation retained at that level.

Each "seller", be it a business unit, a sales division, a sales office, or a sales person, needs to have control over its own forecast and its own commitments. Similarly, it should be allocated its own ATP, over which it exercises full control. It can allocate that ATP down to its member sellers, or retain it at the "seller" level. And it can use the ATP to make promises to customer requests.

In this way, the various sellers form a seller hierarchy. Forecasts and commitments are aggregated up that seller hierarchy, and allocated ATP is disaggregated down that hierarchy.

ATP-by-Product-by-Seller

It is quite common for different sellers to have some control over their own price lists. Thus there is often a need to have all the per-product-characteristic ATP splits for each seller in the hierarchy.

For example, Jill may commit to selling 500 widgets and 100 gadgets while Joe commits to only 50 widgets, but 1000 gadgets. It is important that they both can forecast, commit, and be allocated each of those products separately. They may forecast and therefore be allocated very different mixes.

Due to the desire to sell widgets before gadgets, each seller may be allocated their full widget commitments, but often shorted on their gadget commitments. In the above example, Jill may be allocated 500 widgets and only 10 gadgets, and Joe is allocated 50 widgets and only 100 gadgets.

Alternatively, the sales office may allocate gadgets proportional to the widget allocation rather than the gadget commitment, in which case Jill may get 500 widgets and all 100 gadgets, whereas Joe gets 50 widgets and a mere 10 gadgets.

Note further that product allocations within each seller are not limited to by-item. Each seller could forecast, commit to, and be allocated for different

widget grades, different gadget delivery lead times, different quantity breaks, and so on.

Giving the Customer Options

In the dynamic world of sales, ATP allocations must be flexible. For example, consider a customer who wants a gadget in 6 weeks. There may be no ATP left for 6-week lead time, but some may remain for 2-week lead time. It is not acceptable to turn away the business if the customer says they are willing to pay the higher price for the 2-week product. The lead-time of 2 weeks is a minimum, not a maximum. Thus, creating divisions by just holding back allocations until within the 2-week lead time is unacceptable. Similarly, creating divisions by defining a different item number for each product is unsatisfactory.

If the customer is not willing to pay that higher price, the next question is "When can you get it to me?" Given various ATP products, there may be a whole list of possibilities.

Consider a customer request for 50 slow widgets to be delivered on May 1. The answers may include the following options:

- Get 50 fast widgets on May 1, but you must pay more
- Get 15 slow widgets on May 1
- Get 35 additional slow widgets on May 10

The customer has several choices:

- Get 50 of the fast widgets
- Get 15 of the slow widgets and only 35 of the fast widgets, reducing the cost
- Take 15 of the slow widgets, and the other 35 later
- Wait until May 10 to get all 50 slow widgets
- Take the 15 slow widgets and look to a competitor for the rest
- Reject all the options and go elsewhere

The customer is given the ability to choose among all the possibilities. In this way, the separate ATP allocations can represent different opportunities to the seller and customer rather than rigid restrictions. The options can be presented immediately (on the phone, in an EDI transaction, through an Internet Web Browser, etc.). The resultant promises are completely reliable since the ATP numbers were based on a reliable plan.

5. Cooperatively Integrating Humans with Automated Planning

The Problem-Oriented Planning paradigm provides both automated and manual planning in a way that is consistent with the way human planners want to work. By having the automated algorithms follow the same paradigm, the manual and automated planning efforts become freely intermixable and cooperative.

The Strategy-Driven Planning mechanism allows human planners to effectively direct the automated algorithms to work cooperatively with manual planning efforts. The focus is not on creating a magically wonderful autonomous planning algorithm; rather, the goal is to use the automated planner's extreme speed to automate planning activity that the human planner could do, if he had the time and the inclination. Superior plans come out of the higher speed and more detailed analysis possible by the computer, coupled with a high-degree of cooperation with the human planner, who has been freed of all repetitive work, leaving time for innovative planning. (Of course, Strategy-Driven Planning does not throw out sophisticated optimizing algorithms; it just adapts them to cooperate with the human planner.)

Truly Integrated Planning (TIP) requires integration across several dimensions of the planning process. (See the section "Reengineering the Manufacturing Planning Process" for a more detailed discussion of TIP). This section focuses on integration across two of those dimensions: cooperative integration of manual and automated planning, and integration across successive planning efforts of the on-going planning process.

The Problem-Oriented Planning paradigm was designed to provide sophisticated automated planning and optimization in such a way that it is highly understandable by human planners. Although intelligent optimization should be expected from an advanced planning tool, human interaction with the plan is also needed. Human planners should understand what the optimizer is doing, and the optimizer should be designed to cooperate with a human planner.

The Strategy-Driven Planning mechanism allows human planners to effectively drive the automated algorithms to work cooperatively with manual planning efforts. The human planner can instruct the automated planning algorithms with what is desired, and the automated algorithms go off and take care of it.

Automated Planning Algorithms

There are numerous different planning algorithms employed in modern manufacturing planning software packages. There are simulation-based algorithms, expert systems, constraint engines, linear programming-based algorithms, heuristic algorithms, and genetic algorithms. The relative merits of the various algorithms have been studied, reported, and argued. For the most part, no algorithm has proven to optimize clearly better than the other top algorithms.

Which of the algorithms above typically finds the more desirable plans is debatable. But, it may also be *relatively* unimportant. Although most scheduling researchers and software vendors may enjoy debating which algorithms build the best plans, in most real factories, that is not the real problem that needs to be solved.

The Real Problem

Most software algorithm evaluations involve taking in a model to plan and then generating a plan. Best plan wins. (Or fastest to find best plan wins.) But the real-world doesn't work that way. Usually there is already a plan in place and some changes that must be incorporated. Then it is not just as simple as pressing a button. Instead, the human planners interact with the software and the automated algorithms to modify the plan. How good the plan is after pushing the automated planning button is not the key issue. The key issue is how good is the plan after the human planner and the automated algorithms have both worked on the plan. The evaluation of an automated algorithm should not center on how well the algorithm operates on its own, but rather on how well it works as a team member with the human planners.

Every Issue is Not Modeled

It is impractical to assume that the planning algorithm can always know everything. Most real factory planning involves issues that are not modeled, options that are not normally considered, and trade-offs that cannot be preprogrammed.

If everything is not modeled, then the planning algorithm cannot possibly be expected to find the optimal plan. It does not know all the issues that should be considered when evaluating the plan. And it does not know all the options that are possible for improving the plan. At best it can compute the optimal for what is modeled.

For example, the planning algorithm may know the cost of being late by one day and by two. But does it know the cost of being late for a particular customer who recently got poor service, came down hard on the VP, and threatened to switch to a competitor if such service continued?

It may know the efficiency if one person operates the machine and the efficiency if two people operate it. But does it know that the machine can operate at twice the speed if you put Sally on it? Does it even know that Sally can be put on it (she normally works in a different area)?

Although planning software should be *capable* of modeling everything, it should not assume it *is* modeling everything. It is only modeling the portion that it has been told to model. Human planners do not want to put in every last detail, many of which are only used in extremely rare cases.

It is unreasonable to assume that all information can be made available to the planning software. There are always more possibilities in the real world than the software will be told about. Many of the hard decisions and trade-offs must be made by the human planners.

Planning software should expect the need for human interaction. But most automated planning algorithms are not designed well for human interaction.

Planning is Rarely Automatic

Although a good automated planning algorithm can effectively resolve the vast majority of planning problems, it is impractical to assume that it can solve all planning problems. Even if everything were modeled (which it never is), there are often planning problems that require innovative solutions from the human planner.

On the other hand, for many manufacturers planning is a huge task that could easily consume an army of human planners, unless the majority of it is computer automated.

The automated algorithms should be designed to work in *cooperation* with manual planning efforts, not in lieu of them. Both automated and manual planning are necessary and must be cooperatively integrated. However, the focus of most algorithms is on computing the “optimal” plan most quickly with no interaction with humans, rather than on cooperatively working with the human planners.

Planning is an On-Going Problem

Each time manufacturing planning is performed, the slate is not clean, although much planning software seems to be written assuming it is. The human planners worked very hard the day before to develop a good feasible plan for the next few weeks or months. After a day (or shift or hour) goes by, the plan is still largely valid. It makes no sense to throw out the existing plan and start again. However, many planning algorithms do just that.

Planning is *not* the process of building a plan. It is the on-going process of maintaining a plan. The planning problem rarely involves computing a plan from scratch. It should be a long-term integrated effort, not a series of disjointed repetitions.

Given that yesterday’s plan was formed by both automated and manual decisions and that today’s plan should build from that, the automated algorithm must be designed to incorporate those past manual decisions.

The World is Changing

Continuous improvement—a mantra of the 1980s and 1990s. Rapidly changing customer demands. Agile manufacturing. The only constant is change.

The world is changing and the planning model must change with it. A planning tool should be a competitive weapon, not a hindrance. If competition dictates changes, the planning system should not stand in the way. However, many “sophisticated” planning algorithms are built on models that are not well understood by the human planners; models that they cannot effectively modify. Some require “tuning” to get good results, and altering the tuned instrument requires “re-tuning”.

Very often the planners are unable to effectively modify their planning model. Given that the plans from the planning tool drive the floor, making the change on the floor can render the plans invalid.

In a world where change is the only constant, planning algorithms must be designed to be not only transparent to but also fully modifiable by the human planners. The planners must not need to bring in the software vendor or their IS organization.

Practice Makes Perfect

In any complex activity, practical experience is invaluable. As planners repeatedly plan their world they find patterns and develop techniques for effectively dealing with the problems that come up. Such learning should be effectively captured in the planning tools. The algorithms and rules should be designed to evolve as experience yields improvements.

Because the world keeps changing, this process of improvement never ends. As the world changes, new patterns are found and new techniques are developed.

The “Optimal” Folly

The very notion of an “optimal” plan is dubious. What is “optimal” for this plant? How are late orders factored in? Is two days late worse than one day late? Is a late big order worse than a late small order? What’s worse, a one-day-late big order or a two-day-late small order? Should overtime be run to avoid some late orders? Should lots of extra inventory be carried?

Does every last detail have to be modeled to compute optimal? Yes, but worse: the trade-offs between each thing modeled must be given. That usually means converting every issue to a “cost” or “goodness factor” such that they can be mathematically weighted into one number.

Is optimal even needed? What’s it worth to obtain the optimal plan? If the optimal plan is known, can it be followed? Rarely. There are too many things that cannot be known in advance. Even if everything is modeled, it is not possible to accurately model what cannot be known in advance.

Tossing yesterday’s plan to the wind to switch to another that is ‘more optimal’ is usually a bad idea. Organizations have a certain level of inertia: sales people develop marketing plans, distributors contract carriers, foremen organize people and materials, etc. The disruption of the current plan should be factored into the definition of “optimal”. In other words, the cost of “changing” from yesterday’s plans must be provided.

Which is worse, an additional late order or the complete disruption of today's and tomorrow's plans? Can a mathematical formula for that be given in advance?

Finding the optimal plan is not particularly important to most manufacturing planners. Developing a plan that is feasible (can actually be performed), meets certain critical criteria (e.g., order due dates), and is reasonably good for a few other criteria is all that matters. There are a huge number of plans that are reasonably close to optimal, and any one would be acceptable. After all, the real-world rarely follows the plan exactly anyway. Even if the optimal were found, it probably would not be executed exactly.

The real problem is the on-going process of maintaining a feasible plan that is reasonably near optimal. Solving the problem requires cooperation between automated and human planners.

Cooperating with Human Planners

Many planning algorithms are designed to read in the factory model, status, and demand, and then use some planning parameters or rules provided by the human planner to compute the "optimal" plan. After such a plan is created, the human planner evaluates the result, identifies the problems, and diagnoses the planning parameters that need to be revised to produce better results. The plan is then recreated, with the hope that the new result is superior. That process continues until the human planner is willing to make all the remaining changes manually. This approach is not particularly productive, nor is it intuitive for human planners.

A more natural approach for the human planner is to simply jump in and give the automated algorithms some help by directly manipulating the plan itself. When the planner identifies unsatisfactory plans generated by the automated algorithms, it is most straight-forward to just fix those problems manually and then continue the automated algorithms. This is more natural for a human planner than trying to figure out the planning parameters to change such that the planning algorithm will make the desired changes. If the algorithms indeed need adjusting, then they should be adjusted. But it is unreasonable to assume that all information can be made available to the planning software. It is unreasonable to assume that human planners do not need to get directly involved.

That more natural approach only works, however, if the automated algorithms do not undo the manual edits. The planning algorithms must take the current plan, and the edits performed upon it, as an additional basic input. The desired algorithm is one that can continue to improve the existing plan, primarily adjusting for things that have changed, without squashing the manual edits the human planner worked so hard on.

This is the same requirement as that imposed by the on-going planning process. Yesterday's planning effort should not be thrown away. The desired algorithm is one that can take yesterday's plan and anything that has changed since then, and then produce a new plan without getting rid of the manual edits from the day before.

In fact, automated algorithms can be extremely useful during and after manual editing. Manual editing can generate more new problems than it

fixed, but those new problems may be easier for the computer to solve. The algorithms must be designed for such usage. Many algorithms simply cannot be applied in such a manner. Designing an algorithm that can be used cooperatively with manual editing has a fundamental effect on the nature of the algorithm.

Problem-Oriented Planning

The Problem-Oriented Planning (POP) paradigm was designed to allow manual and automated planning to be freely intermixed. It was designed to be transparent and easily understood by the human planners so that they are comfortable modifying it and working cooperatively with it.

Automate Manual Planning

Although it uses technology from more exotic algorithms (e.g., simulated annealing and heuristic search), the POP algorithm itself is founded on how human planners work most effectively. This is not done because it is necessarily the best way for the computer to work on its own, but because it is the best way for the computer to work in cooperation with humans.

Maskell discusses:

Mathematical models of business problems tend to be complex. Issues dealt with are complex and the mathematical techniques are sophisticated. The result is that the people using the analysis frequently do not understand where the numbers have come from and therefore are unable to validate information or apply their own judgment to the task. In turn, users either unthinkingly follow the computer's lead or manually override it -- thereby making many of the analysis figures useless. People working this way often feel alienated because the only person who can understand the system is some mathematical boffin [expert] working in the MIS department. This situation invariably leads to conflict.¹⁴

Maskell argues:

For this to work, the systems must be simple and transparent to those using them, and these people must be trained to use them effectively. In general, the operations-research concepts are not compatible with this new approach to the management of operations and clerical personnel.¹⁵

The POP algorithm starts with a list of the "problems":

- Plans that cannot or should not be left as they are
- Criteria that are not being fulfilled
- Rules that are not being satisfied

¹⁴B. H. Maskell, *Software and the Agile Manufacturer*, 1994, pp. 14.

¹⁵B. H. Maskell, *Software and the Agile Manufacturer*, 1994, pp. 15.

The algorithm then explores various ways to edit the plan to alleviate those problems. As it makes changes to the plan, it evaluates the overall quality of the altered plan. Note that this is the way human planners work naturally -- they attack the problems (the exceptions).

This paradigm is the same whether automated or manual planning is being performed, making the control of and results from the automated algorithms understandable and highly useful to the human planners. Note that the edits the automated algorithm makes are the same edits that the human planner can make. The list of Problems and the analyses that the automated algorithm uses are also available to the human planner.

The POP automated planner is designed to be a member of the team -- a very fast member (makes thousands of edits and performs thousands of analyses each second) -- but a member of the *team*. The human planners should see the algorithm as a workhorse that they can send off to do exactly what they would do if they had the time (and inclination) to do it themselves. They should not be asking the algorithm to "work magic". They should fully understand exactly what it is doing. Such understanding is fundamental to productive teamwork.

Identifying Problems

Problem-Oriented Planning software continually monitors for planning problems. For instance, all conditions that violate feasibility are problems, such as over-utilized resources, negative balances in inventories, or any other hard-constraint violations. Conditions that affect the desirability of a plan are also identified as problems, such as late orders, short orders, higher inventories, overtime, or higher costs. These problems are identified, collected, and classified according to type and severity. Every time the plan is adjusted, the lists of problems are adjusted as well.

These lists can be filtered based upon type, area, severity, or tolerances, allowing the human planner to obtain the list of problems that are most critical. With each planning decision the planner can immediately see any new problems that are created.

POP automated algorithms are designed to operate on the same paradigm. They can see the problems which must be addressed, know various ways to adjust the plan to eliminate those problems, perform various analyses in order to choose the best adjustments, and then try (and retry repeatedly) various adjustments until a good solution is found.

In either case, the prerequisite to devising great plans is understanding all the problems. If problems go unnoticed, the plans are not optimal. And if those unnoticed problems are feasibility problems, the plans are not even executable.

Allowing Infeasibility

Although feasibility in the final plan is a must, it should not be a hard constraint. Allowing infeasible plans to exist allows more seamless integration between manual and automated algorithms, simplifies many manual planning efforts, and is fundamental to a unique approach to generating optimized plans.

Allowing infeasible plans greatly simplifies manual interaction. For example, if there is some external reason that an operation must be run at noon on Friday (its customer will be touring through at that time), then the human planner would like to be able to simply move that one operation. That may cause the plan to be infeasible. Requiring the planner to simultaneously adjust for the feasibility issues in order to make the move is unreasonable. Far more helpful software allows the planner to create the infeasible state and then invoke the automated algorithms to take care of the infeasibility.

Similarly, a human planner may want to invoke the automated algorithms to generate the best possible schedule, allowing no more than 10% over-utilization of resources. That intermediate plan would be infeasible, but can be quite useful. The human planner can then decide whether to increase capacity (overtime, higher costs) in certain areas. Being able to see which resources could use overtime can be very informative. After making those adjustments, the planner may want to continue the planning algorithms, this time allowing no over-utilization.

Generating and working with infeasible plans can be very useful. The key, of course, is that every infeasibility must be identified and understood. An infeasible plan should only exist temporarily during the planning effort. Such plans should not be released or followed.

Reasonably Near Optimal

Although designed to support seamless intermixing of manual and automated planning, the POP automated planner must also be extremely effective at generating good plans (reasonably near optimal). In that quest, the POP algorithm takes advantage of the fact that infeasibilities are allowed.

For the algorithm, enforcing feasibility at all times simply cuts off valuable parts of the search space. It creates discontinuities that tend to trap the algorithm into local optima. Evaluation of infeasible states can be very informative to the automated planner, just as it is useful to human planners.

Cooperating with Intelligent Local Solvers

The Problem-Oriented Planning paradigm makes it easier for the automated algorithm and human experts to cooperate. It also makes it easier for multiple otherwise independent automated solvers to cooperate.

For example, consider a highly sophisticated expert system for sequencing operations on a resource that has widely varying sequence-dependent setups. Finding a reasonable sequence can be critical and difficult. Now consider upstream from that resource, another resource which must run very large lots and uses a special form of block planning optimization. Neither optimizer tries to incorporate the issues of the other. But a reasonable plan cannot be formed without using the two highly intelligent optimizers. Yet the goal of the overall planning activity is to produce a plan that is globally good.

In the same way that the Problem-Oriented Planning algorithm can incorporate and evaluate decisions made by the automated planner and by several human planners, it can also incorporate decisions made by highly specialized automated planning algorithms. These specialized solvers make

local decisions, and those local decisions are coordinated by the global POP algorithm.

The local solvers must be designed in such a way that they can cooperate with the POP algorithm and thus other local solvers. Still they must remain local solvers (as opposed to growing into complex global solvers).

Cooperating with Human Planners

In the real world, where the human planner must do the “driving”, it is far more important for the results of the planning algorithms to be understandable and usable than optimal. The human planners must be able to build on and work with the planning algorithms.

The automated planning algorithms are easily understood because they are based on the same Problem-Oriented Planning paradigm that the human planners use in manual planning. The algorithms address the same problems that the planner sees, and they use the same plan manipulations to adjust the plan that are available manually. Nothing mysterious is ever performed on the plan. Everything the automated algorithms do could be done manually, albeit much more slowly. The focus of the automated algorithms is not to produce more optimal results than is possible manually, it is to produce the same results that an intelligent human planner would (given adequate time), but much more quickly and easily. Using automated algorithms to take care of numerous, repetitive planning adjustments is completely natural for a human planner.

Not only is the list of remaining problems displayed for the human planner, the analyses used by the automated planning algorithms are made available to the human planner. The same numbers that help the automated algorithms make decisions can be extremely helpful to the human planner. This is because they are both trying to do similar things. The planners must come up with innovative solutions to the problems, but they need not be manually computing ratios, summaries, and other analyses. That is the chore of the computer and the planning software.

Finally, and possibly most importantly, the automated algorithms do not recompute the plan from scratch. The algorithms simply attack the remaining problems. Previous planning decisions are not arbitrarily undone. The planner can freely alternate between making manual adjustments and invoking automated algorithms.

With the ability to freely intermix manual and automated actions comes the desire to use the automated algorithms to perform highly specific tasks, effectively handling some piece of the problem. For that, a mechanism is needed for “driving” the automated algorithms.

Strategy-Driven Planning

Strategy-Driven Planning is a mechanism for controlling the Problem-Oriented Planning algorithm. A “strategy” is a way of directing the automated planner in what to do. It specifies which problems to work on, the goals to achieve, the adjustments that are allowed, and the condition on which to quit.

Individual strategies can be constructed, named, and saved for future use. In this way, organizations can create a library of shared strategies that are useful for solving certain problems.

Ask Your Teammate for Help

As stated earlier, the automated planner should be very much like an additional human planner on the team that just happens to be extremely fast. The automated planner is probably not as creative nor as innovative, but it is very intelligent because it can effectively master all the rules and techniques that the human planners can describe. Coupled with the extreme speed, the automated planner can accomplish things that would take years for a team of human planners to accomplish.

However, the automated planner cannot replace the human planners. The human planners teach the automated planner everything that is reasonably common, but they do not cover all the exceptional issues that can arise. Innovation and creativity remain the domain of human experts.

The human planners are running the show while the automated planner does most of the work. To facilitate that, there must be a way for the human planners to clearly direct the automated planner on what to do. The human planners develop planning strategies for dealing with various situations in their environment. Those strategies must be communicated to the automated planner.

Specify the Problems to Work On

The POP automated planner essentially takes the list of problems and edits the plan to eliminate those problems (if it can). The problems that the algorithms try to eliminate are specified by the strategy being used. Some examples of problems that can be addressed include the following:

- Over-utilized resources
- Heavily utilized resources
- Under-utilized resources
- Inventory below safety stock
- Inventory below zero
- Excess inventory
- Late orders
- Early orders
- High costs

Tolerances can be provided for each problem, allowing small problems to be ignored.

The strategy can also specify the domain to address. The domain indicates which orders, resources, or inventories in which problems should be addressed. The strategy can attack one resource, all the resources and

inventories at a location (work center, cell, department, plant, or site), or all resources.

Note that the manual method of doing this is using the list of problems, selecting the next most interesting problem, and studying it to find a planning adjustment that eliminates the problem.

Specify What's 'Optimal'

While trying to resolve the identified Problems, the automated planner tries to find plans that are the most desirable (most "optimal"). But what is desirable? Optimal?

What is more important: to minimize the number of late orders, or to minimize the total days late of the orders? To minimize inventory or carrying costs? To maximize profit or return on investment? Profit or market share? Profit or customer service?

Different organizations have different goals. Those goals often evolve over time. Or the goals vary depending upon the portion of the overall supply chain that is being planned.

Specify Allowed Adjustments

When attempting to eliminate a problem, there are many planning adjustments that can be used. For instance, an operation can be planned earlier or later, the quantity can be reduced, an alternate routing, operation, or resource may be chosen, operations may be combined, split, or overlapped, orders may be split, and so on.

At any particular time, the human planner may want to restrict the automated algorithms from using certain adjustments. For instance, the planner may want the automated algorithms to try to balance a resource, but without making anything late or short. The strategy specifies which adjustments are allowed. It may be that the goals cannot be achieved with the adjustments allowed. In that case the algorithms simply give up and let the planner decide what to do next.

Specify When to Quit

Finally, a Strategy needs to specify when to stop. That may be when a plan close enough to optimal has been found, or when that strategy has gone long enough and it's time for another.

Termination can be based on the 'optimal' measure, based on the problem list, based on the number of things that have been tried, or based on the time of day. For example, the strategy may specify that the planning algorithm stop when the problems are resolved, or when no problems are solved for 3 minutes, or when a total of 15 minutes has passed.

Multi-Phase Strategies

Sequences of strategies may be linked to form a multi-phase strategy, then named and saved. Each of the strategies that make up a multi-phase strategy are executed one after the other. "Manual" phases can also be added, which

essentially pause the automated activity for the human planner to consider performing some manual planning actions.

In this way, organizations can actually compose their “standard” planning effort into a recorded script. Different users may develop their own multi-phase strategies that they find effective in different situations. They may share their strategies with other planners, who may use or modify them to take care of different issues. Mutual education and superior cooperation between planners is made easier by such sharing of strategies.

The multi-phase strategies can also be used to take advantage of characteristics of the plant. For instance, many Process Flow Scheduling manufacturers are more effective if they plan the various “stages” of their operation in a certain order¹⁶. Each of the stages can be addressed by a strategy phase, and the multi-phase strategy includes the proper ordering.

Strategy-Driven Planning gives the human planner a high level of control over the automated algorithms. The high level flow of the planning effort is effectively programmable via a mechanism that simply describes what the planner would like to achieve and by what means it can be achieved.

Manual Strategies for Guidance

Strategies can be useful for manual planning as well as automated. The human planner appreciates the same information that the automated planner needs: a sorted list of the problems to focus on, a summary of the important planning measurements, and the current value of how optimal the plan is.

Manual strategies can be put into a sequence of strategies so that the automated planner stops at predefined points and asks for human intervention. For example, the automated planner may run 3 or 4 strategies that balance load on all but 3 resources where running overtime is a possibility. That may not be a decision the automated planner is allowed to make, so a manual strategy is inserted to allow a human planner to make that decision *before* proceeding to solve the remaining capacity problems.

In this way, a strategy can encode a full human planning strategy. Rather than writing on a sheet of paper something such as “first balance X, then adjust Y, then check for Z, then...”, those steps can be recorded in strategies. The strategies not only tell what to do next, but automatically compute the list of problems of concern and the criteria for evaluation.

Special Strategies

An “active” strategy is one whose problem list and desirability criteria are maintained with each change to the plan. Both can be displayed at all times while planning activity continues.

The “auto” strategy is an active strategy that is automatically invoked after each planning change. This strategy is usually used to define the problems that the human planners find uninteresting and would like to remain solved at all times. For example, some strategies may have all material problems

¹⁶ S. F. Bolander and S. G. Taylor, “System Frameworks for Process Flow Industries”, *Production and Inventory Management Journal*, Fourth Quarter, 1993, p. 12-17.

“auto” resolved, leaving (or moving) all problems to the resources or orders. This needs to be a fast strategy since it is used after every change.

The “background” strategy is an active strategy that is automatically invoked when the planning system is inactive for some period of time. The strategy can work through millions of alternatives in the search for a superior plan whenever the tool would otherwise be inactive. (Users may find that the plan seems to be better after lunch and greatly improves overnight.)

Some Examples

The following are two familiar and divergent examples that hopefully are useful in illustrating the possibilities for SDP. While these examples are oversimplified to some degree they are useful in many real-world scenarios.

MRP and then Capacity

The MRP (materials requirements planning) algorithm propagates demand upstream through intermediate materials back to raw materials and purchases. It ignores capacity constraints while doing so. The result is an infeasible plan -- but not an uninteresting plan.

The MRP plan tells you what your plan looks like if you have plenty of capacity and can purchase whatever you need. In that sense, it is somewhat of an idealistic plan.

If capacity is pliable (running overtime or hiring temps is easy) starting with an MRP plan can be quite valuable. You can then verify purchases, and if those pass you can adjust capacity to match the plan.

You can direct similar activity by using a strategy. Define a strategy that ignores capacity problems and purchasing problems, but resolves all material problems by requesting more supply (propagates backward or upstream). That gives you an MRP-like plan.

Now to step beyond MRP, the next strategy may attempt to resolve the purchasing problems. Wherever a purchasing problem cannot be resolved upstream, it causes delays downstream. By also continuing to resolve the material problems, but now by delaying or trying alternates, the purchase problems are either resolved via alternates or pushed out to the demand.

This, though infeasible, is also an interesting plan. It tells you what you can do considering purchasing constraints, but ignoring capacity constraints.

The next strategy in this sequence may be a manual one. It may give the human planner a list of all the overloaded capacity problems. The human planner analyzes capacity reports and decides where to run overtime or hire temps. (Alternatively, rules could be set up for each resource that can run overtime to decide whether to do so. This strategy can then be an automated one that makes these decisions.)

With that done, the last strategy resolves the remaining capacity problems by juggling the operation plans to find an ordering that satisfies all the capacity and material constraints.

Process Flow Scheduling (PFS)

As mentioned earlier, Process Flow Scheduling (PFS) is an alternative approach to planning a factory. It breaks the factory flows up into process stages. Each process stage is decoupled from the others by inventory.

Planning starts with the most critical stage and balances capacity there. That leaves material problems at the upstream and/or downstream decoupling inventories. Planning then proceeds to the next stage upstream or downstream. It resolves the material problems just created and the capacity problems in that next stage, leaving material problems at the next decoupling inventory.

This proceeds in both directions. The key is that the material and capacity problems are dealt with one stage at a time, and that there is a preferred order in which the stages are addressed. That preferred order can be encoded in a multi-phase strategy. Each of the sub-strategies resolves the material and capacity problems of one stage only. The multi-phase strategy executes these one after another in the desired order.

Conclusion

Although MRP and PFS are often considered opposing approaches to planning, they do not require different planning software. Both can both be accommodated by a Strategy-Driven Planning tool. That allows for *mixtures* of MRP and PFS in a factory that has elements that need each approach. Such flexibility is a necessity for supply chain planning software since most supply chains contain a wide mixture of different manufacturing types.

Note also that the approaches are encoded easily into strategies. No software development expertise is needed, just planning expertise. Planners can encode their approach to planning their supply chains into strategies. As they grow in expertise they can evolve and refine those strategies. When following the strategy does not resolve all of the problems, the human planner can easily step in and make the innovative planning decisions and then let the strategies proceed. There is no need to adjust the strategies just because exceptional conditions are not covered.

